

BrightStor[®] CA-ASTEX[™] **Performance**

DASD Manager User Guide **2.8**



Computer Associates[™]

This documentation and related computer software program (hereinafter referred to as the "Documentation") is for the end user's informational purposes only and is subject to change or withdrawal by Computer Associates International, Inc. ("CA") at any time.

This documentation may not be copied, transferred, reproduced, disclosed or duplicated, in whole or in part, without the prior written consent of CA. This documentation is proprietary information of CA and protected by the copyright laws of the United States and international treaties.

Notwithstanding the foregoing, licensed users may print a reasonable number of copies of this documentation for their own internal use, provided that all CA copyright notices and legends are affixed to each reproduced copy. Only authorized employees, consultants, or agents of the user who are bound by the confidentiality provisions of the license for the software are permitted to have access to such copies.

This right to print copies is limited to the period during which the license for the product remains in full force and effect. Should the license terminate for any reason, it shall be the user's responsibility to return to CA the reproduced copies or to certify to CA that same have been destroyed.

To the extent permitted by applicable law, CA provides this documentation "as is" without warranty of any kind, including without limitation, any implied warranties of merchantability, fitness for a particular purpose or noninfringement. In no event will CA be liable to the end user or any third party for any loss or damage, direct or indirect, from the use of this documentation, including without limitation, lost profits, business interruption, goodwill, or lost data, even if CA is expressly advised of such loss or damage.

The use of any product referenced in this documentation and this documentation is governed by the end user's applicable license agreement.

The manufacturer of this documentation is Computer Associates International, Inc.

Provided with "Restricted Rights" as set forth in 48 C.F.R. Section 12.212, 48 C.F.R. Sections 52.227-19(c)(1) and (2) or DFARS Section 252.227-7013(c)(1)(ii) or applicable successor provisions.

© 2001 Computer Associates International, Inc.

All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Contents

Chapter 1: Introduction

CA-ASTEX and the DASD Manager Component	1-1
How You Can Use the DASD Manager	1-2
DASD Manager Online Help	1-3
Types of Data the DASD Manager Analyzes	1-4
How CA-ASTEX Uses Interval Data	1-4
Creating a Report Using Data from Multiple Systems	1-4
How CA-ASTEX Uses Trend Data	1-5
CA-ASTEX Documentation	1-5
Contacting Technical Support	1-5

Chapter 2: Using the Allocation Director

What Is the Allocation Director?	2-1
How the Allocation Director Works	2-1
How the Allocation Director Fits Into the HSM Non-SMS Allocation Process	2-2
Finding Out-of-Balance Conditions in SMS	2-4
Data Analyzed by the Allocation Director versus the SRM	2-4
Allocation Director Benefits	2-5
How You Can Use the Allocation Director	2-6
Software Compatibility Issues	2-6

Chapter 3: Displaying Storage Statistics

What Types of Screens Can You View?	3-1
Example of a Storage Statistics Screen	3-2
Procedure for Using the Online Storage Statistics Facility	3-3
Using the Online Help	3-4
Manipulating a Screen	3-5

Accessing Other CA-ASTEX Facilities	3-6
Accessing the Storage Performance Expert	3-6
Accessing the Online Cache Statistics	3-7
Accessing the Command Interface	3-7
Accessing Complex-level Data	3-7

Chapter 4: The Batch Storage Statistics

What Types of Batch Reports Can You Produce?	4-1
How the Batch Reports Differ from the Data Screens	4-2
Procedure for Requesting Batch Reports	4-3
Preparing for a Batch Report	4-3
Using the Online Help	4-4
Looking at Some Sample Selections	4-5
Example 1	4-5
Example 2	4-6
Example 3	4-6

Chapter 5: Interpreting Problem Analysis Data

Understanding the Column Headings	5-1
---	-----

Chapter 6: Interpreting Response Analysis Data

Understanding the Column Headings	6-1
---	-----

Chapter 7: Interpreting Path Analysis Data

Understanding the Column Headings	7-1
---	-----

Chapter 8: Interpreting Cache-RAID Analysis Data

Understanding the Column Headings	8-1
---	-----

Chapter 9: Interpreting SMS Problem Analysis Data

What Is the SMS Problem Analysis Screen?	9-1
Understanding the Column Headings	9-2

Chapter 10: Interpreting Reserve-Allocation Data

Understanding the Column Headings	10-1
---	------

Chapter 11: Interpreting Seek Analysis Data

How to Access the Seek Analysis Screen	11-1
Understanding the Column Headings	11-1

Chapter 12: Displaying the System Summary Screen

What Is the System Summary Screen?	12-1
How to Access the System Summary Screen	12-2

Chapter 13: Displaying the Data Set Analysis Screen

What Is the Data Set Analysis Screen?	13-1
How to Access the Data Set Analysis Screen	13-2
Understanding The Column Headings	13-2

Chapter 14: Displaying the PAV Analysis Screen

What is the PAV Analysis Screen?	14-1
How to Access the PAV Analysis Screen	14-1
Understanding the Column Headings	14-2

Chapter 15: Reorganizing Volumes and Data Sets

Understanding the Reorganization Utilities	15-1
Reorganizing Volumes	15-2
Using FDR/COMPAKTOR	15-3
Using DFDSS	15-3
Unmovable Data Sets	15-5
How the ASXSEEKA Program Works with DFDSS	15-5
Reorganizing Members of a PDS	15-6
Reorganizing Non-partitioned Data Sets	15-7

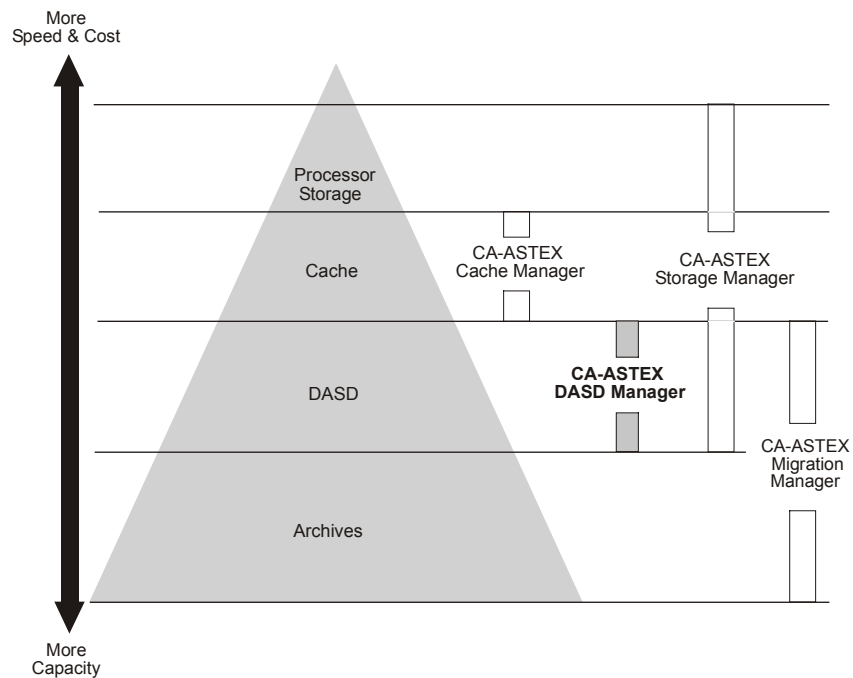
Appendix A: Storage Statistics Display Commands

Command Quick Reference	A-1
-------------------------------	-----

Index

CA-ASTEX and the DASD Manager Component

The DASD Manager component of CA-ASTEX monitors DASD usage and provides detailed statistics and utilities to help you manage your DASD resources. DASD Manager manages the DASD portion of the storage hierarchy, as shown below.



How You Can Use the DASD Manager

When using the DASD Manager component, you work with the three DASD Manager facilities (shown highlighted below) and the Allocation Director feature. To select one of the facilities, you enter its selection number (for example, D1 for the Online Storage Statistics) at the option line.

```
CA-Automated Storage Expert 2.8
Primary Menu

OPTION ==>

Storage Manager                                DASD Manager
S1 Storage Performance Expert                 D1 Online Storage Statistics
S2 Configuration Manager                     D2 Batch Storage Statistics
S3 Workload Reporter                         D3 Reorganization Utilities

Cache Manager                                CA-ASTEX Facilities
C1 Cache Management Expert                   A1 User Defaults
C2 Online Cache Statistics                   A2 Database Utilities
C3 Batch Cache Statistics                     A3 Command Interface
C4 Batch Cache Simulation

Migration Manager
M1 Online Migration Statistics
M2 Batch Statistics
```

Online Storage
Statistics

Let you examine detailed DASD I/O measurements for the current interval (real-time) or a specified previous interval. The data is presented on a series of screens that you can view using ISPF.

Batch Storage
Statistics

Let you examine detailed DASD I/O measurements for previous intervals via a series of offline reports.

The Reorganization
Utilities

Provide ways to make improvements to DASD performance. These utilities let you reorganize volumes, reorganize members of a partitioned data set, and analyze access patterns for data sets.

The Allocation
Director

The Allocation Director is a “behind-the-scenes” feature of CA-ASTEX that *automatically* directs allocations away from problem volumes and paths. The Allocation Director is described in more detail in the chapter "Using the Allocation Director".

DASD Manager Online Help

Use CA-ASTEX's ISPF-based help system to guide you as you use the DASD Manager.

We recommend that you review the "Table of Contents" help screen for each facility first. "Table of Contents" screens show you all of the topics related to a particular facility. The following screen is an example "Table of Contents" for the DASD Manager component:

```

Tutorial-----Tutorial
Option ==>
      *-----*
      |          DASD Manager - Table of Contents          |
      *-----*
The DASD Manager component provides detailed DASD statistics and
reorganization utilities. The DASD Manager consists of the three
facilities listed below. For more detailed information on any one
facility, enter its selection code in the Option field.

Option          Explanation
1 Online Storage Statistics  Let you examine detailed DASD I/O
                             measurements for the current interval
                             or a specified previous interval.
2 Batch Storage Statistics   Let you produce reports that show
                             detailed DASD I/O measurements for
                             previous intervals.
3 Reorganization Utilities   Let you reorganize volumes and data
                             sets.

                        UP= Main Help Screen  F3= Exit Help
  
```

You can access this screen in two ways:

- At the CA-ASTEX Primary Menu, move your cursor to the desired facility and press **F1**.
- From any help screen within the DASD Manager component, type UP at the command line.

Types of Data the DASD Manager Analyzes

DASD Manager analyzes trend and interval data collected by CA-ASTEX common measurement routines.

How CA-ASTEX Uses Interval Data

When using the Online Storage Statistics facility, you can view single intervals of data (current or previous intervals) from a single system, or you can use the CA-ASTEX File Merge/Conversion utility to:

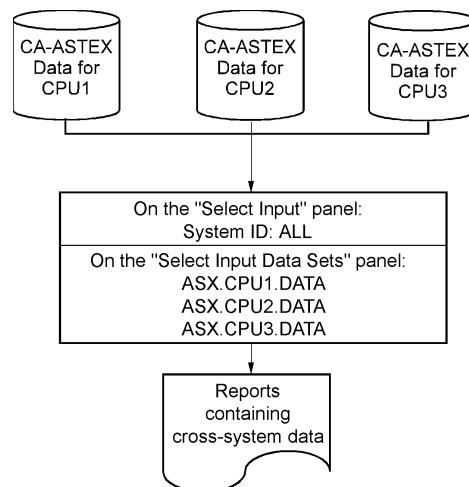
- Combine multiple intervals into one larger interval
- Combine interval data from multiple systems into one database

See the *CA-ASTEX User Guide* for more information on the CA-ASTEX File Merge/Conversion utility.

When using the Batch Storage Statistics facility, you can run reports for a single interval (previous intervals only) or combined intervals, and for a single system or multiple systems.

Creating a Report Using Data from Multiple Systems

The following illustration shows how performance data from multiple systems is incorporated into a single report using the Batch Storage Statistics facility:



How CA-ASTEX Uses Trend Data

The Allocation Director uses data from the CA-ASTEX Trend Database to make its allocation decisions. This trend data provides long-term information on paths and volumes.

CA-ASTEX Documentation

The following manuals are included with the CA-ASTEX product:

Getting Started

Administrator Guide

User Guide

Storage Manager User Guide

Cache Manager User Guide

DASD Manager User Guide

Migration Manager User Guide

Message Guide

Contacting Technical Support

For technical assistance with this product, please contact Computer Associates Technical Support at esupport.ca.com. Technical support is available 24 hours a day, 7 days a week.

Before contacting technical support, please try to resolve your problem by reading this manual and related publications such as the “Contacting Technical Support” appendix of the CA-ASTEX *Administrator Guide*, and by using any Help systems associated with the product.

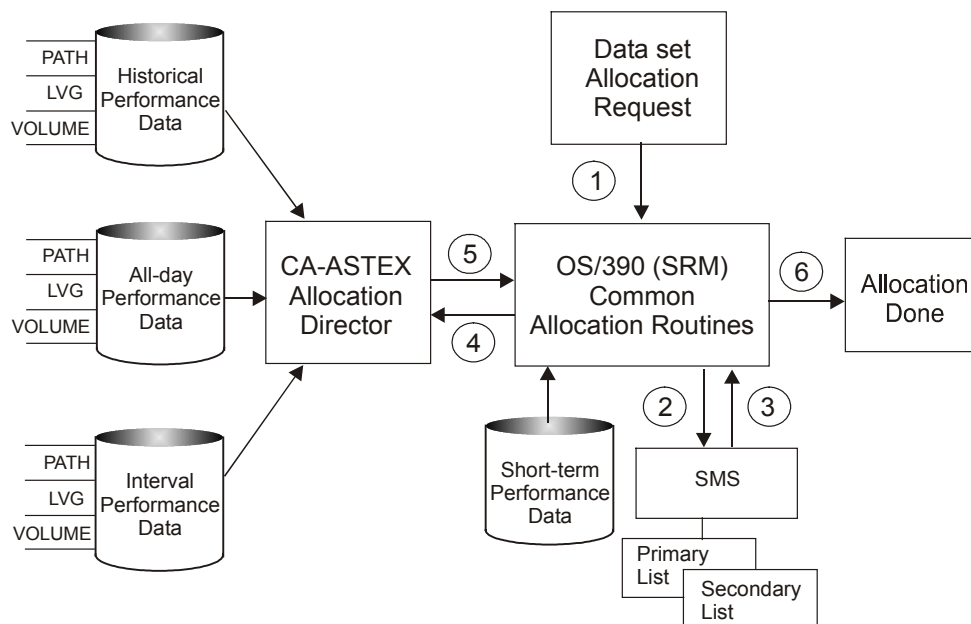
Using the Allocation Director

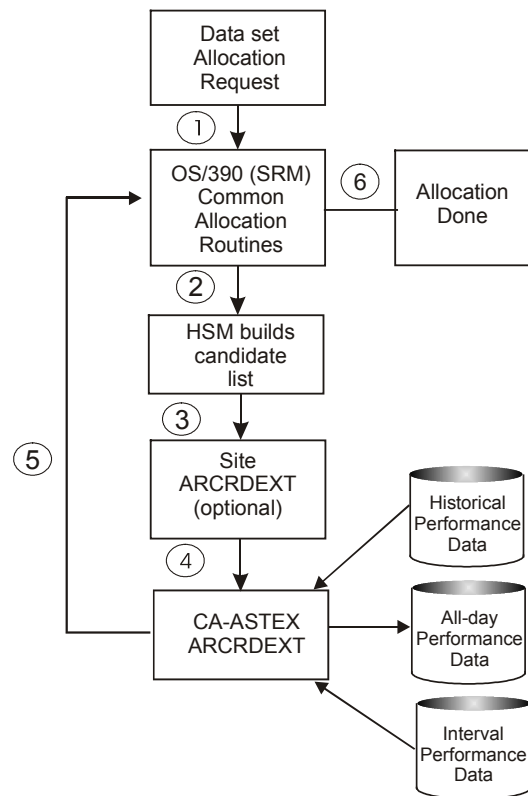
What Is the Allocation Director?

The Allocation Director is a feature of the DASD Manager that automatically directs allocations away from problem volumes and paths. The Allocation Director works “behind-the-scenes,” with little intervention from you.

How the Allocation Director Works

The Allocation Director works in conjunction with the OS/390 common allocation process and SMS. It uses CA-ASTEX all-day, interval, and trend (historical) data to make the best allocation decisions, as shown below:





How the Allocation Director Fits Into the HSM Non-SMS Allocation Process

The Allocation Director improves upon the way OS/390 handles new allocations and allocations due to HSM-recalled data sets. For new non-SMS allocations, common allocation passes a list of candidate volumes to the System Resources Manager (SRM). For new SMS allocations, the SMS allocation routines build a primary list of candidate volumes and pass that list to SRM. Using short-term volume and path utilization data from the past interval, the SRM chooses a device for allocation.

For SMS-managed data sets being recalled by HSM, control is passed to the SMS allocation routines. For non-SMS managed data sets, HSM builds a candidate list of volumes based on recall attributes and space criteria. If your site has a user-written HSM recall exit, ARCRDTEXT, the candidate list is then passed to this exit. This exit can be used to re-order the candidate list based on the criteria you have defined. Common allocation then passes the volume on the top of the list. If allocation fails for that volume, the next volume is passed until the list is exhausted.

In contrast, the Allocation Director makes the best allocation decisions *before* a poor volume is selected. The Allocation Director uses CA-ASTEX trend, all-day, and interval data to select the best volume from the candidate list. The Allocation Director also ensures that the volume is in the allocation list from which the SRM made its choice. For non-SMS managed recalled data sets, the Allocation Director ensures that the volume selected is in the candidate list that HSM built (or, if you have one, in the candidate list built by your site's ARCRDTEXT exit). CA-ASTEX provides its own ARCRDTEXT to optimize any non-directed DFHSM recalls of data sets that are not managed by SMS.

As a result, the Allocation Director can make allocation decisions based on a volume's long-term performance, taking into account such factors as:

- Path delays
- ESCON delays
- Logical Volume Group delays
- Volume queuing
- Seek time
- Cross-system and dispatching delays
- Volume and path utilization
- Cache utilization (back-end staging)

For example, if a volume currently has low utilization, the SRM might label the volume as a good allocation candidate. The SRM would then select the volume for an allocation. The Allocation Director, on the other hand, looks at the volume's long-term performance and may find that it has consistently missed response time objectives over the past few weeks, thus making it a potentially poor candidate. The Allocation Director would then avoid this volume.

Once the Allocation Director finds a set of good candidates, it balances allocations across those candidates so that no one volume is overloaded. The Allocation Director periodically rebuilds its allocation tables to take performance changes into account.

Finding Out-of-Balance Conditions in SMS

The Allocation Director can help you to determine when SMS is producing an *out-of-balance* condition (not choosing the best allocation candidates). This condition can lead to poor resource utilization.

During allocation processing, SMS separates allocation candidates into two lists:

1. **The Primary List** contains volumes that meet or exceed service level and space objectives.
2. **The Secondary List** contains volumes that meet space objectives only.

The primary list is passed to the SRM; the SRM then chooses a volume for allocation. In this case, the Allocation Director can work with the SRM to ensure that the best allocation decisions are made.

If all of the candidates are in the secondary list, SMS bypasses the SRM call and the lists remain unbalanced. SMS then selects the volume with the greatest free space from the secondary list. Although the Allocation Director cannot help balance performance in this case, it does provide allocation counts on the DASD Manager "Reserve-Allocation" screens and reports. See the chapter "Interpreting Reserve-Allocation Data" for a complete description of these screens and reports. These counts can inform you of when and how often this out-of-balance condition is occurring. You can then take the appropriate action to resolve the problem.

Data Analyzed by the Allocation Director versus the SRM

The following table contrasts the types of data the Allocation Director and the SRM use to make allocation decisions:

Types of Data	Allocation Director	SRM
Trend data	Yes	No
All-day data	Yes	No
Interval and real-time data	Yes	Yes
Device and path utilization statistics	Yes	Yes
Device and path problem counts	Yes	No
LVG utilization statistics	Yes	No
LVG problem counts	Yes	No
Cache back-end staging statistics	Yes	No

Allocation Director Benefits

Using the Allocation Director offers the following benefits:

Continuous monitoring of OS/390 allocation	The Allocation Director continuously examines OS/390 allocation decisions. If the Allocation Director finds a better device than the one that OS/390 found, that device is chosen for allocation.
Proactive solutions to allocation problems	The Allocation Director examines performance data over a wider time span than OS/390 does. The Allocation Director's data starts at the current interval and goes back several weeks. As a result, the Allocation Director's decisions are <i>well informed</i> . The best volumes are chosen to avoid data set moves and volume swaps that are necessary to recover from a poor allocation decision.
No adverse effect on good allocation candidates	If the Allocation Director finds that OS/390 has chosen the best candidate for allocation, it does not interfere with the normal OS/390 allocation process.
Prevent allocations from "clumping" on the best device	The Allocation Director balances allocations among a group of best candidates so that no one device is overloaded. Also, the Allocation Director continually adjusts to changes in workload by rebuilding its list of best volumes in each storage group.
Complements SMS Allocation decisions	The Allocation Director assures the best volume is selected from the primary list, and informs you if SMS is using the secondary list volumes to perform allocation.
Complements HSM Recall allocations	The Allocation Director assures the best volume is selected for SMS and non-SMS recalled data sets and informs you of the number of allocations due to recalling data sets.

How You Can Use the Allocation Director

Follow these steps to use the Allocation Director:

Step	Action
1 (Optional)	Install the CA-ASTEX ARCRDEXT exit. See the CA-ASTEX <i>Administrator Guide</i> for instructions on how to install the exit. Also, see the description of the ALLOCHSM parameter.
2	Initialize the Allocation Director. Refer to the discussion of the ALLOCINT parameter in the CA-ASTEX <i>Administrator Guide</i> .
3	Activate the Allocation Director for the desired storage group(s). Refer to the discussions of the ALLOCDIR parameter in the CA-ASTEX <i>Administrator Guide</i> .
4	Update the Trend Data Base daily (preferably around midnight). See the CA-ASTEX <i>User Guide</i> for detailed instructions on how to update your Trend Data Base.

Software Compatibility Issues

In addition to the Allocation Director, other vendors' software products interface with common allocation and the SRM. POOL-DASD from Empact Software provides an SRM driver that actually replaces the SRM device allocation function. This replacement driver uses a device allocation algorithm that is based solely on space criteria, whereas the Allocation Director uses performance information for device selection.

These are two completely different approaches to DASD management. You should analyze your needs to determine which approach is best for you.

Displaying Storage Statistics

What Types of Screens Can You View?

DASD Manager's Online Storage Statistics are organized hierarchically so you can examine intervals of data at various levels of detail. The Online Storage Statistics are for the knowledgeable performance analyst who wants to examine detailed DASD I/O measurements.

Data is presented in seven types of screens:

Screen	Description
Problem Analysis	Shows how current measurements compare with your predefined response objectives. Points out specific performance problems and lets you dynamically access the Storage Manager for discussions of these problems and their solutions.
Response Analysis	Shows average I/O response time measurements and their components.
Path Analysis	Shows path utilization statistics.
Cache-RAID Analysis	Shows cache utilization statistics. Also shows information associated with Physical Disk Contention (PDC) problems due to RAID 1, RAID 5, and RAID 6 technology.
SMS Problem Analysis	Shows I/O performance measurements against SMS storage class objectives.
Reserve-Allocation	Shows reserve and allocation statistics for individual volumes and groups of volumes in a shared DASD environment.
Seek Analysis	Shows seek statistics for either data sets or jobs.

The data on each screen is organized according to a contention group “view” or a storage group “view.” And, from each screen, you can analyze data at various levels of detail:

- Complex
- System
- Path
- Group
- Storage Class
- Logical Volume Group (LVG)
- Volume
- Data Set
- Job

Example of a Storage Statistics Screen

Here is a typical Storage Statistics screen with pointers to the key fields:

Diagram labels and pointers:

- Detail field**: Points to the 'Detail' field in the summary line.
- Summary field**: Points to the 'Summary' field in the summary line.
- Command area**: Points to the 'Command' area in the summary line.
- Column identifier line**: Points to the 'Column' identifier line in the summary line.
- Screen field**: Points to the 'Screen' field in the summary line.
- View field**: Points to the 'View' field in the summary line.
- Interval field**: Points to the 'Interval' field in the summary line.
- System ID**: Points to the 'System' field in the summary line.
- Collection time**: Points to the 'Collection' time in the summary line.
- Collection date**: Points to the 'Collection' date in the summary line.
- Summary line**: Points to the 'Summary' line in the table.
- Selection column**: Points to the 'Selection' column in the table.
- Detail lines**: Points to the 'Detail' lines in the table.

Group	Cache Status	Cand I/O%	Tot%	Cand%	R%	RS%	DFW%	CFW%	Response Imp%	Loads Elim%	Busy Imp%
System	95.2	89.7	94.2	96	99	90	0	87.9	0	63.7	
DCG096	CD	99.6	93.4	93.7	97	100	88	0	88.5	0	58.6
DCG240	CD	89.6	85.2	95.0	93	97	99	0	85.6	0	75.9
D00064	CD	93.3	89.7	96.1	95	97	100	0	88.0	0	80.0
DCG000	CD	0.0	0.0	0.0	0	0	0	0	0.0	0	0.0
D00320	CD	0.0	0.0	0.0	0	0	0	0	0.0	0	0.0

Procedure for Using the Online Storage Statistics Facility

Follow these steps to display online storage statistics.

Note: Press **F1** from any CA-ASTEX screen to receive help in completing the following procedure. The online help facility explains each CA-ASTEX screen in detail.

Step	Action
1.	Select Online Storage Statistics (option D1) from the CA-ASTEX Primary Menu.
2.	Select input data by filling in the "Select Input" screen. <ol style="list-style-type: none">(Optionally) By typing YES where it asks "Change the expert analysis defaults?": you can change the expert analysis defaults by filling in the "Select Expert Analysis" screen.(Optionally) By typing YES where it asks "Change the display defaults?:" you can change the display defaults by filling in the "Select Display Defaults" display.
3.	Select a display from the "Display Selection List" screen.
4.	View the data.

Using the Online Help

Use CA-ASTEX's ISPF-based help system to guide you as you use the Online Storage Statistics.

We recommend that you first review the "Table of Contents" help screen for the Online Storage Statistics. The "Table of Contents" screen shows you all of the topics related to the Online Storage Statistics facility.

```
Tutorial-----Tutorial
Option ==>
      *-----*
      | Online Storage Statistics - Table of Contents |
      *-----*
The Online Storage Statistics facility lets you look at detailed
performance data. To view a specific topic about the Online Storage
Statistics facility, enter its selection code in the Option field.

1  Selecting Input Data                13 Command Interface Screen
2  Selecting an Interval               14 Problem Analysis Screen
3  Selecting Input Data Sets           15 Response Analysis Screen
4  Selecting Expert Analysis Defaults  16 Path Analysis Screen
5  Selecting Display Defaults          17 Cache-RAID Analysis Screen
6  Understanding the Display Format     18 SMS Problem Analysis Screen
7  Moving Between Screens              19 Reserve-Allocation Screen
8  Selecting a Level of Detail         20 Seek Analysis Screen
9  Switching Between Views             21 Data Set Analysis Screen
10 Printing a Display                 22 System Summary Screen
11 Issuing Display Commands           23 Function Key Settings
12 Accessing Other Components

      UP= Help Menu  F3= Exit Help
```

You can access this screen in two ways:

- At the CA-ASTEX Primary Menu, move your cursor to the Online Storage Statistics facility and press **F1**.
- From any Online Storage Statistics help screen, type UP at the command line.

Manipulating a Screen

Once you reach the initial screen, you can begin manipulating it in any of the following ways:

- Move to other types of screens (for example, the Path Analysis, Reserve-Allocation, Seek Analysis).
- View data at various summary and detail levels.
- Switch between “views” of contention and storage group data.
- Print a screen.
- Issue display commands to manipulate displayed data.

See the Online Storage Statistics “Table of Contents” help screen for detailed information on how to manipulate a screen.

Accessing Other CA-ASTEX Facilities

From any Storage Statistics screen, you can access the following CA-ASTEX facilities:

- The Storage Performance Expert, where you can read the problem and solution discussions
- The Online Cache Statistics facility, where you can review detailed cache statistics
- The Command Interface, where you can issue CA-ASTEX operating commands

Accessing the Storage Performance Expert

If you type a P (for “Problem List”) in the selection column to the left of one of the lines of data, the Storage Performance Expert displays a problem list that contains all the problems found at that detail level and below.

For example, if you type P to the left of a path name, as shown below, the Storage Performance Expert lists all problems on that path; if you type P to the left of a data set name, the Storage Performance Expert lists all problems pertaining to that data set. From the problem list you can then display the problem discussion or solution discussion.

Press **F3** to return to the data screen from the problem list.

Summary: System		View: Contention		Interval 9:00 - 10:00 as of 9:56					
Detail: Group		Screen: Path Analysis		System: IPOX Date: 24 JAN 01					
>				Scroll HALF					
-									
Group	I/O %	--Connect-- Avg	RPS %	RPS Avg	ESCON DPB	-Pth Sys%	Prb%- Grp%	-Prb RPS%	Dist- ESC%
.a.....c.....u.....f1.....g1.....z.....h1.....i1.....b1.....c1.....									
System	100.0	6.4	100.0	5.1	1.1	4.9		65	35
P PATH06	29.5	10.0	37.2	10.3	1.3	2.1	9.0	79	21
PATH05	29.3	6.1	38.5	2.1	2.6	0.6	1.5	47	53
PATH01	21.7	3.5	7.9	2.7	1.0	0.2	1.3	67	33
PATH04	10.3	6.9	10.7	6.3	0.8	1.2	10.9	62	38
PATH02	9.2	3.5	5.7	6.5	0.2	0.7	6.4	23	77

Note: The “P” command is valid only if the Storage Manager component is authorized and you requested that expert analysis be performed.

Accessing the Online Cache Statistics

If you type CSTATS on the command line, you see an Online Cache Statistics screen. The screen is at the highest level of detail. To return to the Online Storage Statistics, type DSTATS. If you exit the Online Cache Statistics, you return to the DASD Manager's "Select Input" screen or the "Intervals Available" screen.

Accessing the Command Interface

The Command Interface is a CA-ASTEX facility where you can issue CA-ASTEX operator commands without leaving the Online Storage Statistics screens.

To access the Command Interface, enter the CMD command from the command line of any Storage Statistics screen.

At the "Input: Modify CA-ASTEX" prompt, you can specify any of the same operands you would include on one of the F ASX operator commands that are described in the *CA-ASTEX User Guide*.

To leave the Command Interface, press **F3/15**. You return to the screen where you issued CMD.

Accessing Complex-level Data

If you are using a merged Interval Database as input, you can view data for your entire storage complex.

An example of a complex-level Response Analysis screen is shown below.

You can view data for a particular system in the complex by typing an S to the left of a system name, as shown below.

```
Summary: Complex View : Contention Interval 15:00 - 15:30 as of 15:30
Detail : System Screen: Response Analysis System: ALL Date: 26 JAN 01
>
      I/O      -Prob%-  Rsp  ----Response Components----- RPS  ESCON
      System % Mode Cmp% Avg IOSQ Pend Conn Disc Dsp Avg DPB
.....a.....C.....p.....e.....q.....S.....t.....u.....v.....w.....g1.....z...
Complex 100.0      9.7      8  0.1  0.5  4.5  1.3  2.3  0.2  0.0
PROD    36.2      5.5      9  0.1  0.7  5.0  1.4  2.3  0.5  0.0
S S008   26.7      2.5      6  0.0  0.4  3.3  0.2  2.1  0.0  0.0
S999    19.1      0.9      8  0.0  0.2  3.3  1.0  3.4  0.1  0.0
S028    10.0      0.3      9  0.3  0.9  4.0  1.2  3.0  0.3  0.0
TEST     6.3      0.2     14  0.0  0.2  7.0  5.6  0.7  0.0  0.0
S032     1.7      0.3     17  0.0  0.2 11.8  1.4  3.2  0.6  0.0
SYS1     0.0      0.0      0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
SYS2     0.0      0.0      0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
***** BOTTOM OF DATA *****
```

You can also view group- and volume-level data for the complex by typing a G or a V at the Selection column. An example of group-level data is shown below.

```

Summary: System View : Contention      Interval 10:30 - 11:00 as of 10:32      ROW 1 TO 6 OF 6
Detail : Group  Screen: Problem Analysis  System: S008      Date: 29 AUG 01
>
                                     Scroll HALF
Group      Total  I/O  I/O Problem  --Prob%--- --Problem Distribution %---
          I/O    %   SEC  Type    Sys%  Grp% Pth PDC Xsy Vol  Sk Dsp Dsn
....a.....b.....c.....d.....e.....f.....g..j3...h....i....j...f2...k.
System      7,712 100.0  47.6      D    2.8      8 11  0  <1  3  8 69
PATHEMC6    4,172  54.1  25.8      D    2.1   3.9 10  0  0  1  4  4 82
PATHEMC1    1,311  17.0   8.1      0.2   1.4  0  0  0  0  0 17 83
PATH02      1,031  13.4   6.4      0.1   0.4  0  0  0  0 25 50 25
DCG009       682   8.8   4.2      0.4   4.5  0 81  0  0  0 19  0
PATHEMC4     516   6.7   3.2     <0.1   0.6  0  0  0  0  0  0 100
PATHRAM1       0   0.0   0.0      0.0   0.0  0  0  0  0  0  0  0

```

The Batch Storage Statistics

What Types of Batch Reports Can You Produce?

The Batch Storage Statistics are organized hierarchically so you can examine intervals of data at various levels of detail. The Batch Storage Statistics show you detailed DASD I/O measurements in the form of easy-to-use batch reports.

The Batch Storage Statistics provide the following types of printed reports:

Report	Description
Response and Exception Analysis	Shows average I/O response time measurements and their components. Also shows how current measurements compare with your predefined response objectives.
Path Analysis	Shows path utilization statistics.
Cache-RAID Analysis	Shows cache utilization statistics. Also shows the number of Physical Disk Contention (PDC) problems associated with RAID 1, RAID 5, and RAID 6 devices.
SMS Problem Analysis	Shows I/O performance measurements against SMS storage class objectives.
Reserve-Allocation	Shows reserve and allocation statistics for individual volumes and groups of volumes in a shared DASD environment.
Seek Analysis	Shows seek statistics for data sets.

The reports have the following structure:

- An optional overview section that contains total performance statistics for the period of time encompassed by the report.
- Separate sections that divide the report period into smaller intervals of time. Each section shows performance data collected only during that interval of time.

The data on each type of report is organized according to a contention group or storage group “view.” And, for each type of analysis, you can show data at various levels of detail:

- Complex
- System
- Path
- Group
- Storage Class
- Volume
- Data set
- Job

How the Batch Reports Differ from the Data Screens

Essentially, the Storage Statistics facilities present the same performance data in both online and batch forms. However, there are a few differences between the online and batch facilities and their output:

- Both the screens and reports use the CA-ASTEX Interval Database as input. The screens can also use *real-time* data, which is data from the current interval. The batch reports can use SMF data as input.
- The batch facility can automatically merge intervals and data from multiple systems. If you want to view data for multiple intervals or systems via the online facility, you must first merge intervals or databases using the CA-ASTEX merge utility.
- The reports and screens contain the same columns of data, with two exceptions:
 1. The screens provide a “Problem Analysis” and a “Response Analysis” screen, whereas the batch facility combines the data on those screens into one “Response and Exception Analysis” report.

Note: An “exception” is a condition that occurs when an I/O request exceeds your response time objective.

2. Reports that have detail lines containing job or data set data show an “Active Time” column of data in place of the “Last Access” data that you find on data screens using real-time data. The “Active Time” measurements show the total amount of time each data set or job was active during the period covered by the report.

Procedure for Requesting Batch Reports

Follow these steps to produce a Batch Storage Statistics report.

Note: Press **F1** from any CA-ASTEX screen to receive help in completing the following procedure. The online help facility explains each CA-ASTEX screen in detail.

Step	Action
1	Prepare for the Batch Report. See “ Preparing for a Batch Report ” in this chapter for tips on how to prepare for a batch report.
2	Select Batch Storage Statistics (option D2) from the CA-ASTEX Primary Menu.
3	Specify the scope and content of the report by filling in the “Select Input” panel. (Optional) You can edit the JCL before submitting the job by typing YES.
4	Select the type of report by filling in the “Select Reports” panel.
5	Select your input data sets by filling in the “Select Input Data Sets” panel.

Preparing for a Batch Report

You may want to check the values of several parameters in the ASXPARM member of the ASX.CNTL data set. These parameters determine the amount of data that appears on the report:

- Values on the IDB, VOLSMFT, DSNSMFT, and JOBSMFT parameters that enable CA-ASTEX to collect performance data
- An interval length on the DINTVL parameter
- An EM or DM value on the IMODE parameter
- A response time objective on the MSR parameters. This value determines how much data set data is collected when CA-ASTEX is running in exception mode (EM), and it also defines what CA-ASTEX considers to be a problem.

Note: An “exception” is a condition that occurs when an I/O request exceeds your response time objective.

Using the Online Help

Use the CA-ASTEX ISPF-based help system to guide you as you use the Batch Storage Statistics.

We recommend that you first review the “Table of Contents” help screen for the Batch Storage Statistics. The “Table of Contents” screen shows you all of the topics related to the Batch Storage Statistics facility.

```
Tutorial-----Tutorial
Option ==>
      *-----*
      | Batch Storage Statistics - Table of Contents |
      *-----*
The Batch Storage Statistics facility lets you produce batch reports
of detailed CA-ASTEX performance data.

To view a specific topic about the Batch Storage Statistics facility,
enter its selection code in the Option field.

1 Selecting Input Data for the Batch Reports
2 Selecting the Type of Report You Want
3 Selecting an Input Data Set

UP= Help Menu  F3= Exit Help
```

You can access this screen in two ways:

- At the CA-ASTEX Primary Menu, move your cursor to the Batch Storage Statistics facility and press **F1**.
- From any Batch Storage Statistics help screen, type UP at the command line.

Looking at Some Sample Selections

The following examples show how the criteria you specify on the two batch report panels controls report content.

Example 1

Suppose that you specified the following criteria on the “Select Input” panel:

```

Start date:      01 JAN 2001
Stop date:       01 JAN 2001
Start time:      00 : 00
Stop time:       24 : 00

What length should the interval reports be?      : 08 : 00
View the data by CONTENTION ... STORAGE group?   : STORAGE
Report only on this group:    IMS
  
```

Then, on the report list screen, suppose that you selected the following report, with all of its summary/detail combinations:

Response-Exception		
Summary	Detail	Analysis
System	Group	S_
Group	Class	S_
Group	Volume	S_
Group	DSN	S_
Group	Job	S_
Volume	Class	S_
Volume	DSN	S_
Volume	Job	S_
Class	DSN by Group	S_
Class	DSN by Volume	S_

Using all this criteria, and other defaults, the batch routines produce a “Response and Exception Analysis Report” containing performance information for a storage group named “IMS.” The report contains an overview section and three other sections: one containing data from 12:00 a.m. to 8:00 a.m., a second one for 8:00 a.m. to 4:00 p.m., and the third for 4:00 p.m. to 12:00 a.m. Each section of the report shows all valid “Summary/Detail” combinations of system, group, volume, class, data set, and job data from all systems at your site.

Example 2

Suppose that you specified the following criteria on the “Select Input” panel:

Start date:	01 JAN 2001
Stop date:	01 JAN 2001
Start time:	00 : 00
Stop time:	24 : 00
What length should the interval reports be:	: 01 : 00
View the data by CONTENTION ... STORAGE group?	: CONTENTION
Report only on this group:	ALL

Then, on the report list screen, suppose that you selected just one summary/detail combination for the following report:

	Path
Summary Detail	Analysis
Group DSN	S_

Using all this criteria, and other defaults, the batch routines produce a “Path Analysis Report” containing performance information for all contention groups that are defined on your system. The report contains 24 sections, each one containing data for a one-hour interval of time (12:00 a.m. to 1:00 a.m., 1:00 a.m. to 2:00 a.m., and so on). Each section of the report shows “Group” data on the “summary line” and “Dataset” data on the “detail lines.”

Example 3

Suppose that you specified the following criteria on the “Select Input” panel:

Start date:	01 JUL 2001
Stop date:	01 JUL 2001
Start time:	09 : 00
Stop time:	15 : 00
What length should the interval reports be?	: 00 : 20
View the data by CONTENTION ... STORAGE group?	: STORAGE
Report only on this group: CICS*	Generics combined: NO

Then, on the report list screen, suppose that you selected one summary/detail combination for the following two types of reports:

Path		
Analysis		
Summary Detail		
Volume	DSN	2_
Group	Class	—
Group	Volume	—
.	.	
.	.	
.	.	
Seek		
Analysis		
Summary Detail		
Volume	DSN	1_
Class	DSN by Volume	—

Using all this criteria, and other defaults, the batch routines produce a “Seek Analysis Report” followed by a “Path Analysis Report.” The reports contain performance information for each volume in any storage groups that have the letters “CICS” at the beginning of their names. Each report contains 18 sections, and each section contains data for a twenty-minute interval of time (9:00 a.m. to 9:20 a.m., 9:20 a.m. to 9:40 a.m., and so on). Both reports show “Volume” data on the “summary line” and “Dataset” data on the “detail lines.”

Interpreting Problem Analysis Data

Understanding the Column Headings

The Problem Analysis screens and reports identify the existence of path, logical volume group (LVG), volume, and data set problems. Use these screens to find where your response objectives are not being met.

This chapter describes the columns of data found on the “Problem Analysis” screens. The following information is provided (where applicable) for each column:

- Description of the column or group of columns
- How DASD Manager calculates the value
- How you can use the information

A sample Problem Analysis screen is shown below. As you read the description of each column heading, refer to this example.

```

ROW 1 TO 6 OF 6
Summary: System View : Contention Interval 12:00 - 12:30 as of 12:20
Detail : Group Screen: Problem Analysis System: S008 Date: 29 AUG 01
> Scroll HALF

Group      Total I/O  I/O  I/O Problem --Prob%--- --Problem Distribution %---
          I/O    %   SEC Type  Sys% Grp% Pth PDC Xsy Vol Sk Dsp Dsn
....a.....b.....c.....d.....e.....f.....g...j3...h...i...j...f2...k.
System    125,089 100.0 102.4      D  3.3      7  4  0  2  3 10 74
PATHEMC6  45,960  36.7  37.6      D  2.2  5.9  7  0  0  1  3  5 83
PATHEMC4  29,961  24.0  24.5      0.3  1.3  1  0  0  3  1  7 88
DCG009    17,588  14.1  14.4      0.3  2.1  0 45  0  4  0 51  1
PATHEMC1  17,369  13.9  14.2      D  0.2  1.6  9  0  0  1  2 13 75
PATH02    14,211  11.4  11.6      0.3  2.6 16  0  0  4  9  8 63
PATHRAM1      0   0.0   0.0      0.0  0.0  0  0  0  0  0  0  0
***** BOTTOM OF DATA *****

```

Note: Some columns appear on more than one type of screen, so if you cannot find an explanation of a particular column in this chapter, look up the column heading in the index.

Column a

summary/detail values

This column identifies the specific “Summary” and “Detail” data being displayed. In this chapter on the sample [Problem Analysis screen](#), the “summary” data shows system-wide measurements, while the “detail” data shows the performance of four contention groups (PATHEMC6, PATHEMC4, and so on).

The first horizontal row of data conveys information at the “Summary” level. The remaining rows of the screen contain “Detail” level data.

On screens that show detailed volume data, an asterisk (*) after a volser means that CA-ASTEX is currently collecting job and data set data for this volume. A plus sign (+) indicates that exception mode is specified and that job and data set data was collected at some point during the current interval but is not being collected now.

Column x1

(Volser)

The “(Volser)” heading appears only on screens that have a detail field value of Dataset.

It lists the names of volumes where the data sets in Column “a” reside.

Column y1

(Job name)

The “(Job name)” heading appears only on screens where the detail field value is Dataset.

It lists the jobs that were the last to access the corresponding data sets listed in Column “a.”

Column b

Total
I/O

This column shows the total number of I/O requests issued within the entire system (125,089) and for each contention group (45,960 for PATHEMC6, 29,961 for PATHEMC4, and so on).

This is the total number of I/Os that have occurred since 12:00 p.m. (the beginning of the indicated 12:00 to 12:30 p.m. INTERVAL). In this chapter on the sample [Problem Analysis screen](#), the system issued 125,089 I/O requests between the time CA-ASTEX was started (12:00 p.m.) and the time this screen was created (12:20 p.m.).

If the total number of I/Os for the system seems small, you may not be able to draw many valid conclusions from the remaining measurements.

Column c

I/O
%

This column shows what percentage of the system's total number of I/O requests went to each contention group. In this chapter on the sample [Problem Analysis screen](#), the 100.0% at the top of this column stands for the 125,089 I/O requests counted. 36.7% of those requests went to the PATHEMC6 contention group, 24.0% went to the PATHEMC4 group, and so on.

The number is calculated by dividing the total I/O count for each contention group by the total I/O count for the system. (45,960 divided by 125,089 equals 36.7%, and 29,961 divided by 125,089 is 24.0%.)

If your total I/O load is well balanced, all the percentages in this column should be almost equal.

To find what is contributing to an unbalanced I/O load among your contention groups, look at the group that has a higher-than-average percentage.

Column d

I/O
SEC

This column shows the number of I/O requests issued each second.

The number is calculated by dividing the Total I/O number by the number of seconds that have passed during the specified interval. In this chapter on the sample "[Problem Analysis screen](#)," 125,089 I/Os divided by 1,260 seconds (the number of seconds between 12:00 p.m. and 12:20 p.m.) equals 102.4 I/Os per second.

If your I/O load is evenly balanced, the number of I/O requests per second should be roughly the same for each contention group listed. If you want to balance the I/O load among your contention groups, look at the group that has the highest number of I/O requests per second.

P, G, V, and D Flags

Problem Type

The “Problem Type” column contains flags that indicate whether path (“P”), logical volume group (“G”), volume (“V”), or data set (“D”) problems exist.

Problem flags that appear on the system line indicate the kind of problems that exist on the system as a whole.

On a summary line, the problem flags indicate what problems exist at the summary level and below.

On a detail line, the problem flags indicate all problems associated with that detail level.

For example, if “V” and “D” flags appear on a detail line for a storage group, it means that volume and data set problems were found on the paths that are part of the storage group.

If you did not request that expert analysis be performed (via the “Select Defaults” screen), an asterisk (*) appears in place of problem flags on the summary line.

Column e

--Prob%-- Sys%

Note: The abbreviation in this column heading matches the value in the “Summary” field. For a “Summary” value of “System,” the abbreviation is “Sys%.” For “Group,” the abbreviation would be “Grp%,” for “Logical Volume Group,” it would be “LVG%,” for “Volume,” it would be “Vol%.” For “Complex,” the abbreviation would be “Cmp%.”

“Prob%” represents the percentage of I/Os whose response times did not meet your defined response objectives. In this chapter, on the sample “[Problem Analysis screen](#),” 3.3% of the system’s 125,089 I/O requests exceeded response time objectives. 2.2% of those response problems went to the PATH EMC6 contention group, and 0.3% are for the PATH EMC4 group.

The number is calculated by dividing the total number of response problems (for the system and each contention group) by the number at the top of the Total I/O column.

Usually, you should investigate whatever contention group has the highest percentage of response problems. (See the next heading on “Grp” response problems.)

To improve overall system performance, you should move data sets or reschedule jobs away from contention groups that have high response problem percentages.

Column f

--Prob%--
Grp%

Note: The abbreviation in this column heading matches the value in the "Detail" field. For a "Detail" value of "Group," the abbreviation is "Grp%." For "Logical Volume Group," the abbreviation would be "LVG%." For "Volume," the abbreviation would be "Vol%." For "Dataset" or "Job," it would be "Dsn%" or "Job%." For "System," the abbreviation would be "Sys%."

This column shows what percentage of the Total I/O value on each detail line (in this case, for each contention group) were response problems. (5.9% of the PATHEMC6 group's 45,960 I/O requests were response problems, while 1.3% of the PATHEMC4 group's 29,961 I/O requests were problems.)

The number is calculated by dividing the total number of response problems for each contention group by the corresponding I/O total.

Watch this percentage if you are evaluating I/O response times for a particular contention group rather than the entire system. Investigate the contention group that has the highest percentage of response problems.

Column g

--Problem Distribution %--
Pth

This column shows the percentage of the total response problems caused by high RPS time and ESCON delays. See Column g1 in the "Interpreting Path Analysis Data" chapter for an explanation of how average RPS time statistics are calculated.

The number is calculated by dividing the number of path problems by total response problems.

Use the problem distribution percentages to determine in what area you need to take corrective action.

Use the "Path Analysis" screens to determine who the largest contributors to path busy are. Look at the "Path-Prb%" column on the "Path Analysis" screens to determine if your path problem is due to ESCON director port busy or due to RPS delay.

Column j3

--Problem Distribution %---
PDC

This column shows the percentage of the total response problems caused by physical disk contention (PDC) associated with RAID 1, RAID 5, and RAID 6 DASD architectures. High disconnect times for a single volume indicate a PDC problem for a specific volume. High disconnect times for multiple volumes within a logical volume group (LVG) indicate a PDC problem that affects the entire LVG.

PDC problems are caused by contention for a logical volume or group's physical disks. This contention can occur because multiple data sets on a volume or within an LVG are poor cache users and must frequently access the physical disks. Back-end cache staging can also play a role in creating physical disk contention.

The number is calculated by dividing the number of PDC problems by total response problems.

Use the Cache-RAID Analysis screen to determine which data sets are receiving a negative benefit from using cache. Data sets with a negative net benefit should be moved to a non-RAID device and disabled from cache usage.

This column can also display internal RAID contention (IRC) problems associated with RAID 6 technology.

Column h

-- Problem Distribution %---
Xsy

This column shows the percentage of the total response problems caused by high device pending time. This usually indicates a delay because another system in the complex has the device reserved or busy, but it could also indicate delays due to cache back-end staging, channel path, or control unit busy.

The number is calculated by dividing the number of cross-system delay and reserve problems by total response problems.

Use CA-ASTEX to determine if your cross-system delay problems are due to excessive busy conditions or by reserves being issued to other systems.

You should eliminate reserves on volumes whose non-busy reserve times are greatest if there is a reserve problem or move data sets to other volumes to reduce volume busy.

Column i

--Problem Distribution %--
Vol

This column shows the percentage of the total response problems caused by high device IOS queue time. This usually indicates a delay due to excessive device busy.

The number is calculated by dividing the number of volume problems by total response problems.

Use the volume data set screens to determine who the largest contributors to volume busy are.

Data sets need to be moved off the volume. Use the data screens to identify the best data sets to move and where to move them.

Column j

--Problem Distribution %--
Sk

This column shows the percentage of the total response problems caused by excessively high device seek times between different data sets.

The number is calculated by dividing the number of seek problems by total response problems.

Move to a "Seek Analysis" screen to identify the data sets contributing to the high seek problem rate.

You may want to use one of the Reorganization Utilities to perform a seek analysis to help reduce the high seek times.

Column f2

--Problem Distribution %--
Dsp

This column shows the percentage of total response problems caused by high I/O dispatch times.

The number is calculated by dividing the number of dispatch problems by total response problems.

If a high dispatch exception exists, you may wish to investigate changing your weighting for partitions under PR/SM or MDF or your CPU time slicing under VM. Or you may be having an internal IOS processing problem.

Note: If you replied NO to "Should dispatch problems be analyzed?" on the "Select Defaults" panel, the value in this field is zero.

Column k

--Problem Distribution %--
Dsn

This column shows the percentage of response problems caused by I/O to a particular data set. These problems may result from a large number of users trying to share the same data set or by excessively large seek times within the data set (intra-seek problems).

The number is calculated by dividing the number of data set problems (response problems attributed to an individual data set) by total response problems.

Move to the "Data Set Analysis" screen for any data set with a large percentage of response problems.

Move to a "Response Analysis" screen that has a detail value of "Dataset." The data set information on that screen helps you further identify the problem.

Interpreting Response Analysis Data

Understanding the Column Headings

The Response Analysis screens and reports show average I/O response time measurements and their components. You can use these screens and reports to pinpoint the reason why your response objectives are not being met, and the length of the delay.

This chapter describes the columns of data found on the “Response Analysis” screens. The following information is provided (where applicable) for each column:

- Description of the column (or group of columns)
- How DASD Manager calculates the value
- How you can use the information

A sample Response Analysis screen is shown below. As you read the description of each column heading, refer to this example.

```

Row 1 to 9 of 9
Summary: Complex View : Contention Interval 15:00 - 15:30 as of 15:30
Detail : System Screen: Response Analysis System: ALL Date: 26 JAN 01
> Scroll CSR
System      I/O      -Prob%-  Rsp  ----Response Components----- RPS  ESCON
            %      Mode  Cmp%  Avg  IOSQ  Pend  Conn  Disc  Dsp  Avg  DPB
...a.....c.....p.....e.....q.....s.....t.....u.....v.....w.....g1....z...
Complex  100.0      D      9.7      8      0.1      0.5      4.5      1.3      2.3      0.2      0.0
PROD      36.2      D      5.5      9      0.1      0.7      5.0      1.4      2.3      0.5      0.0
S008      26.7      D      2.5      6      0.0      0.4      3.3      0.2      2.1      0.0      0.0
S999      19.1      D      0.9      8      0.0      0.2      3.3      1.0      3.4      0.1      0.0
S028      10.0      D      0.3      9      0.3      0.9      4.0      1.2      3.0      0.3      0.0
TEST       6.3      D      0.2     14      0.0      0.2      7.0      5.6      0.7      0.0      0.0
S032       1.7      D      0.3     17      0.0      0.2     11.8      1.4      3.2      0.6      0.0
SYS1       0.0      D      0.0      0      0.0      0.0      0.0      0.0      0.0      0.0      0.0
SYS2       0.0      D      0.0      0      0.0      0.0      0.0      0.0      0.0      0.0      0.0
***** BOTTOM OF DATA *****

```

Column a

summary/detail values

This column identifies the specific “Summary” and “Detail” data being displayed. The sample “[Response Analysis screen](#)” shows a “Complex/System” screen. It displays the names (SMFIDs) of individual systems (PROD, S008, S999 and so on) belonging to the complex.

The first horizontal row of data conveys information at the “Summary” level. The remaining rows of the screen contain “Detail” level data.

If the Detail Level is Volume, the asterisk (*) after each volume name means that CA-ASTEX is currently collecting job and data set data for this volume. A plus sign (+) indicates that the volume is in exception mode and that job and data set data was collected at some point during the current interval but is not being collected now.

If the “Summary” level is Volume, a **P** after the volser indicates that the volume is defined as a Parallel Access Volume (PAV).

Column x1

(Volser)

This column is the same as the “x1” column on the Problem Analysis Screen.

Column m

DUA

This column contains the device unit address (08C7, 165F, and so on) of each volume.

The “DUA” column heading appears only on screens that have “Detail” values of Volume.

Columns c, e, and f

I/O	-Prob%-
%	Cmp%

Columns “c,” “e,” and “f” are the same columns that appear on the Problem Analysis Screen.

Note: The heading abbreviations in Columns “e” and “f” change with the values in the “Summary” and “Detail” fields. Column “f” is shown only on screens containing detailed group or volume data.

Column p

Mode

The “Mode” heading appears only on screens containing detailed group or volume data. The abbreviations you find in this column describe the current setting of the IMODE parameter for each group or volume listed.

The following is a list of possible abbreviations:

- D** Detail mode is currently specified.
- D-** Mode was changed during this interval. Detail mode is currently specified.
- E** Exception mode is currently specified.
- E-** Mode was changed during this interval. Exception mode is currently specified.
- xxT** Tracing is active for this particular group or volume.
- xxC** The volume has been added or reassigned by the Configuration Manager.
- *** The group or volume has been added by the Configuration Manager, or data is not currently being collected.

Column pa

#PAV

This column shows the number of Parallel Access Volume (PAV exposures for the volume. A null value indicates tat the volume is not defined as a PAV.

This column appears only when the “Detail” level is Volume. See the chapter “Displaying the PAV Analysis Screen” for more information concerning this column heading.

Column q

Rsp
Avg

This column shows the average time (in milliseconds) needed for an I/O request to finish processing. This is the average I/O response time that CA-ASTEX monitors.

The number is calculated by dividing the total amount of time spent processing I/O requests by the total number of I/O requests issued.

The IOSQ, Pend, Conn, Disc, and Dsp averages explained on the next few pages are all components of the response time average. As the sample [Response Analysis screen](#) shows, system S008's IOSQ (0.0), Pend (0.4), Conn (3.3), Disc (0.2), and Dsp (2.1) values add up to an average response time of 6 milliseconds (rounded to the nearest millisecond).

Important! Lowering this average response time for jobs and data sets you care most about is your major objective in DASD tuning.

Column s

---Response Components---
IOSQ

This column shows the average time (in milliseconds) that elapses from the time a user program issues an I/O request until the channel subsystem accepts the "Start Subchannel." As the sample "[Response Analysis screen](#)" shows, the queue time average for the entire complex is 0.1 milliseconds. On system S028, the average queue time is 0.3 milliseconds.

The number is calculated by dividing the sum of all queue times by the total number of I/Os.

At the volume level, a high queue average for a single volume indicates volume contention. High queue averages for all volumes in a contention group indicate a path contention problem.

Contention groups or volumes that have consistently high queue times have too much I/O traffic. Look at the job and data set screens to see which jobs you might reschedule and which data sets you might move to other volumes and/or paths.

Column t

---Response Components---
Pend

This column shows the average amount of time an I/O request waits in the channel subsystem queue before the I/O operation actually starts on the device.

The number is calculated by dividing the total pending time by the total number of I/Os.

Large pending times usually indicate delays due to shared DASD contention, but they could also indicate delay due to cache back-end staging, channel path, or control unit busy.

Look at average RPS time (on the "Path Analysis" screen). If it is also high, that indicates a path problem. If not, the problem is cross system reserves or device busy and CA-ASTEX should be used in other systems to identify the contributor.

Column u

---Response Components---
Conn

This column shows the average amount of time an I/O device spent searching for and actually transferring data. In the sample "[Response Analysis screen](#)," the average connect time across the entire complex is 4.5 milliseconds. For system PROD, the average connect time is 5.0 milliseconds.

The number is calculated by dividing the total connect time by the total number of I/Os.

High connect times indicate problems with long channel programs, long searches, or very large block sizes.

Column v

---Response Components---
Disc

This column shows the average amount of time (in milliseconds) an I/O device was disconnected and not transferring or searching for data. The average disconnect time for the entire complex is 1.3 milliseconds, while for system S028, the average time is 1.2 milliseconds.

The number is calculated by dividing the total disconnect time by the total number of I/O requests.

High disconnect times indicate excessive DASD arm movement (seeks) or high RPS times or ESCON delay.

Look at "Path Analysis" or "Seek Analysis" screens to find a reason for the high disconnect times.

Column w

---Response Components---
Dsp

This column shows the average amount of time (in milliseconds) that an I/O operation waited to be dispatched either by IOS (the I/O subsystem) or the channel subsystem, or for the I/O interrupt to be processed. The sample "[Response Analysis screen](#)" shows that the average delay that the systems in the complex experience is 2.3 milliseconds. I/O requests from system S008 experience somewhat shorter average delays of 2.1 milliseconds.

Note: An asterisk (*) in this column indicates that this value could not be calculated because an insufficient amount of data was collected by the channel subsystem.

This value is calculated by measuring elapsed time from the point when the IOS SSCH (start subchannel) is issued until the interrupt is received back from the channel subsystem. Connect, disconnect, and hardware pending time is subtracted from this elapsed time. The remainder is dispatch time.

This number may be high because of delays in the XA channel subsystem, because of delays in CPU interrupt processing, or because of VM processing (if the OS/390 system is a guest on a VM host system). You may wish to investigate changing your weighting for partitions under PR/SM, your CPU time splicing under VM, or the CPENABLE parameter in the IEAOPTxx member of SYS1.PARMLIB.

Column n1

Last
Access

Note: This heading appears only on “Response Analysis” screens with a “Detail” value of Dataset.

This column lists the time of day, during the current interval, when each data set was last accessed or when a job last accessed a volume.

The most recently accessed jobs and data sets are nearest the top of the screen, and their Last Access times are close to the current time. Suppose your screen has the following heading:

Interval 9:00-10:00
As of 9:56 on . . .

You might see the following times under the Last Access heading:

9:55:02
9:52:49
9:41:58

This tells you when activity for the displayed job or data set was last identified.

Column n2

Active
Time

Note: The “Active time” heading appears only on *batch reports* with a detail field value of Job or Dataset.

This is the total amount of time that each job or data set experienced I/O activity during the period of time covered by the report.

Column g1

RPS
Avg

This column shows the average amount of time an I/O spent waiting due to a rotational positioning sensing miss.

Average RPS time is calculated by dividing TOTAL RPS time by the TOTAL I/O count. RPS time is not measured but is derived using disconnect time, search time and seek time. This derived value may not be 100% accurate in a shared DASD environment due to seeking involving more than one CPU.

Large average RPS times indicate I/Os are suffering long delays due to path busy. If you are experiencing ESCON director port busy (a non-zero value for ESCON DPB), this can in turn cause high RPS times.

In the sample “[Response Analysis screen](#)” the largest average RPS time (0.6 milliseconds) belongs to system S032.

Column z

ESCON
DPB

This column shows the average amount of time an I/O request waits due to ESCON director port busy. High values here indicate that you may be experiencing ESCON director port busy delays.

The number is calculated by dividing the total director port busy time by the total number of I/Os.

Review the “Problem Analysis” screen to see if you are experiencing path problems. If you are, review the “Path Analysis” screen to see if ESCON delay is causing your path problems.

Column z1

Program
Name

Note: The “Program name” heading appears only on a “Response Analysis” screen that has a detail field value of Job or Dataset.

If the detail field value is Dataset, this shows the name of the last program that accessed each data set.

If the detail field value is Job, this shows the name of the last program that the job used.

Interpreting Path Analysis Data

Understanding the Column Headings

The Path Analysis screens and reports show path utilization statistics. You can use these screens and reports to discover the path components that are contributing to poor response time.

This chapter describes the columns of data found on the “Path Analysis” screens and reports. The following information is provided (where applicable) for each column:

- Description of the column or group of columns
- How DASD Manager calculates the value
- How you can use the information

A sample Path Analysis screen is shown below. As you read the description of each column heading, refer to this example.

```

Summary: Group   View: Contention      Row 1 to 4 of 4
Detail: Dataset Screen: Path Analysis  Interval 9:00 - 10:00 as of 9:56
>                                          System: IPOX Date: 24 JAN 01
>                                          Scroll HALF
-

Group      I/O  --Connect-  RPS  ESCON  -Pth Prb%-  -Prb Dist-
          %    Avg    %    Avg    DPB    Lst% Dsn%    RPS% ESC%
x1.y1..a....c.....u....f1....g1....z.....h1..i1.....b1...c1...

C3C0      100.0  6.7  100.0  0.8  0.6    0.2          65  35

SYS808 EIMSVS.RESLIB
TTPX26   11.1   5.3    8.7  0.1  0.2    0.0  0.0    79  21
SYS808 EIMSVS.TFORMAT
TTPX26    6.2   8.7    8.1  0.2  0.2    0.0  0.1    47  53
PRDE08 PI1.DBA00.ARTIKEL.DARTDE
TTPX26    6.1   2.9    2.7  1.5  1.2    0.0  0.6    67  33
PRD809 PV1.WKF50.GENIP.ESDS
MVSDL1    5.9   4.1    3.7  0.1  0.1    0.0  0.0    62  38

```

Columns x1 and y1

(Volser) (Jobname)

These columns are the same as the “x1” and “y1” columns on the Problem Analysis Screen.

Column a

summary/detail values

This column lists the names of individual data sets (EIMSVS.RESLIB, EIMSVS.TFORMAT, and so on) and jobs belonging to the C3C0 contention group.

If the “Detail” value were Volume, this screen would list individual volumes associated with the C3C0 contention group.

If the “Detail” value on this screen were Volume, an asterisk (*) after certain volume names would mean that CA-ASTEX was currently collecting job and data set data for that volume. A plus sign (+) would indicate that the volume is in exception mode and that job and data set data was collected at some point during the current interval but is not being collected now.

If the “Detail” value on this screen was System, and you were using a merged Interval Database as input, this screen would list the systems for your complex.

Column c

I/O
%

This column is the same as the “I/O%” column on the Problem Analysis Screen.

Column u

--Connect--
Avg

This column shows the average “connect time,” which is the average amount of time an I/O device spent searching for and actually transferring data.

Although the column heading is different, the values are the same as those under Column “u” on the Response Analysis Screen.

Column f1

--Connect--
%

This column shows what percentage of the contention group's total connect time went to each data set. In the sample "[Path Analysis screen](#)," 8.7% of the C3C0 group's total connect time went to EIMSVS.RESLIB, 8.1% went to EIMSVS.TFORMAT, and so on.

The number is calculated by dividing the total connect time for each data set by the total connect time for the contention group.

You can use this screen to identify the data sets that contribute the most to "path busy" conditions. In the sample "[Path Analysis screen](#)," data set EIMSVS.RESLIB contributed the most to "path busy" conditions.

If you have a large number of path problems, then move items with the largest values in this column to other paths. If this path is on an ESCON director, and your path problems are due to ESCON delays (see the "Prob Dist ESC%" on this screen), move items with the largest values in this column to other paths not sharing this ESCON director.

Column g1

RPS
Avg

This column shows the average amount of time an I/O spent waiting due to a rotational positioning sensing miss.

Average RPS time is calculated by dividing TOTAL RPS time by the TOTAL I/O count. RPS time is not measured but is derived using disconnect time, search time and seek time. This derived value may not be 100% accurate in a shared DASD environment due to seeking involving more than one CPU.

Large average RPS times indicate I/Os are suffering long delays due to path busy. If you are experiencing ESCON director port busy (a non-zero value for ESCON DPB), this can in turn cause high RPS times.

In the sample "[Path Analysis screen](#)," the largest average RPS time (1.5 milliseconds) belongs to the PI1.DBA00.ARTIKEL.DARTDE data set.

Note: Because most of today's DASD architectures include device (actuator) level buffers at the physical device level, RPS problems due to path busy are extremely rare.

Column z

ESCON
DPB

This column shows the average amount of time an I/O request waits until the ESCON director is free to handle the I/O.

Although the column heading is different, the values are the same as those under Column “z” on the Response Analysis Screen.

Column h1

-Pth Prb%-
Lst%

Note: Abbreviations in this column heading relate to values in the screen’s “Summary” field. In this example, the abbreviation “Lst%” means that the column’s values apply only to the “group” of data sets listed in this screen. They do not reflect measurements for the entire C3C0 contention group.

A path problem is an I/O request that exceeds your response time objective because of high RPS times and/or high ESCON delay times.

This column shows what percentage of the group’s total I/O requests exceeded your response time objective because they had high RPS times and/or high ESCON delay times.

The number is calculated by dividing the total number of path problems (first for the contention group and then for each data set) by the total I/Os for the group.

The data set with the largest “Lst%” value suffers the most from path problems.

To reduce the number of path problems, you need to remove path components that have large “Connect%” values.

Column i1

-Pth Prb%-
Dsn%

Note: The abbreviation in this column heading matches the value on the screen's "Detail" field. In this example, the abbreviation "Dsn%" corresponds to the detail value of Dataset.

This column shows what percentage of the data set's total I/O requests exceeded your response time objective because they had high RPS times and/or high ESCON delay times.

The number is calculated by dividing the total number of data set path problems by the total number of I/Os for the data set.

This percentage indicates how much path delays (due to RPS or ESCON delays) affect overall response time to a particular data set.

To reduce the number of path problems, you need to remove path components that have large "Connect%" values.

Column b1

-Prb Dist-
RPS%

This column shows the percentage of each data set's path problems that consist of RPS delays.

The number is calculated by dividing the total number of RPS problems by the total number of path problems.

To reduce the number of RPS problems, you need to move path components that have large "Connect%" values to other paths.

Column c1

-Prb Dist-
ESC%

This column shows the percentage of each data set's path problems that consist of ESCON director port busy.

The number is calculated by dividing the total number of ESCON problems by the total number of path problems.

To reduce path problems, move path components that have large "Connect%" values to other paths not sharing the same ESCON director(s).

Interpreting Cache-RAID Analysis Data

Understanding the Column Headings

The Cache-RAID Analysis screens and reports show cache utilization statistics and the percentage of problems due to Physical Disk Contention.

This chapter describes the columns of data found on the “Cache-RAID Analysis” screens and reports. The following information is provided (where applicable) for each column:

- Description of the column (or group of columns)
- How DASD Manager calculates the value
- How you can use the information

A sample Cache-RAID Analysis screen is shown below. As you read the description of each column heading, refer to this example.

```

Summary: Group      View : Contention      Interval 11:00 - 11:30 as of 11:19
Detail: Dataset    Screen: Cache-RAID Analysis System: S008      Date: 03 AUG 01
>
Group      I/O    -PDC Prb%- Cand    Wrt -Cache Hits-  --I/O Time Benefit--
%    Lst% Dsn%    %    %    Lst% Dsn%    Saved    Cost    Net
.x1.y1..a.....c.....h3....i3....j1....j4....l1.....m1.....i8.....j8.....k8...
C3C0      100.0    5.0      70.0    3.0  76.6      20m      40m      -20m
SYS808      EIMSVS.RESLIB
TTPX26      46.4    3.7      8.0    50.3    5.8  23.8  53.9      2.0m      1.0m      1.0m
SYS808      EIMSVS.TFORMAT
TTPX26      33.1    1.2      3.7    65.7    12.6  31.1  91.3     500ms     350ms     150ms
PRDE08      PI1.DBA00.ARTIKEL.DARTDE
TTPX26      16.3    0.0      0.0    20.8    0.0  20.2 100.0     100ms     800ms     -700ms
PRD809      PV1.WKF50.GENIP.ESDS
MVSDL1      2.7    <0.1      0.6    33.5 100.0      1.5  95.3     17.0m     35.0m     -18.0m
***** BOTTOM OF DATA *****

```

Columns x1 and y1

(Volser) (Jobname)

These columns are the same as the “x1” and “y1” columns that appear on the Problem Analysis Screen.

Column a

summary/detail values

This column lists the specific “Summary” and “Detail” data being displayed.

The first horizontal row of data conveys information at the “Summary” level. The remaining rows of the screen contain “Detail” level data.

On screens that show detailed volume data, an asterisk (*) after a volser means that CA-ASTEX is currently collecting job and data set data for this volume. A plus sign (+) indicates that exception mode is specified and that job and data set data was collected at some point during the current interval but is not being collected now.

If the “Detail” value on this screen was System, and you were using a merged Interval Database as input, this screen would list the systems for your complex.

Column c

I/O
%

This column is the same as the “I/O%” column on the Problem Analysis Screen.

Column h3

-PDC Prb%-
Lst%

Note: The abbreviations in this column heading relate to values in the screen’s “Summary” field. In this example, the abbreviation “Lst%” means that the column’s values apply only to the “group” of data sets listed in this screen. They do not reflect measurements for the entire C3C0 contention group.

This column shows what percentage of the group’s total I/O requests exceeded your response time objective because of Physical Disk Contention (PDC), which is associated with RAID 1, RAID 5, and RAID 6 technology.

This number is calculated by dividing the total number of PDC problems (first by the contention group and then for each data set) by the total I/Os for the group.

The data set with the largest Lst% value suffers the most from physical disk contention problems.

To reduce the number of PDC problems, workload must be balanced across all available volumes and controllers. If PDC problems are due to disk accesses by poor cache users or cache back-end staging, you need to move busy data sets with low hit percents to other less busy volumes. If PDC problems are occurring because of a high write-to-read ratio, move busy data sets with high DFW counts to other less busy volumes. If PDC problems are occurring at the LVG level, data sets must be moved to volumes that are not part of the problem LVG to correct the problem.

This column can also display internal RAID contention (IRC) problems associated with RAID 6 technology.

Column i3

-PDC Prb%-
Dsn%

Note: The abbreviation in this column heading matches the value in the “Detail” field. In this example, the abbreviation “Dsn%” corresponds to the detail value of Dataset.

This column shows what percentage of the data set’s total I/O count exceeded your response time objective because of Physical Disk Contention (PDC).

This number is calculated by dividing the total number of PDC problems by the total number of I/Os for each detail line.

To reduce the number of PDC problems, the workload must be balanced across all available volumes and controllers. If PDC problems are due to disk accesses by poor cache users or cache back-end staging, you need to move busy data sets with low hit percents to other, less busy volumes. If PDC problems are occurring because of a high write-to-read ratio, move busy data sets with high DFW counts to other, less busy volumes. If PDC problems are occurring at the LVG level, data sets must be moved to volumes that are not part of the problem LVG.

This column can also display internal RAID contention (IRC) problems associated with RAID 6 technology.

Column j1

Cand
%

This column shows the percentage of each data set's total I/Os that are cache candidates. These include reads, sequential reads, and DASD and cache fast writes.

The number is calculated by dividing the data set cache candidate I/O count by the total data set I/O count.

This percentage indicates what percentage of the data set's I/O requests are candidates for a cache hit. Data sets with high cache candidate ratios are good candidates for cache.

Column j4

Wrt
%

This column shows the percentage of I/Os that are write requests.

This number is calculated by dividing the total number of write I/Os (normal writes, DASD fast writes, and cache fast writes) by the total number of I/Os.

Data sets with a low number of writes are possible candidates to move to a non-RAID device. Another option for RAID 6 devices is to rewrite a data set with a low WRT% and a high PDC%. The data is written to a new location, which may resolve the problem.

Note: If you are not authorized to use the Cache Manager Component of CA-ASTEX, "N/A" appears in this column.

Column l1

-Cache Hits-
Lst%

Note: Abbreviations in this column heading relate to values in the screen's "Summary" field. In this example, the abbreviation "Lst%" means that the column's values apply only to the "group" of data sets listed in this screen. They do not reflect measurements for the entire C3C0 contention group.

This column shows what percentage of the group's total cache candidate I/O count consists of read, read sequential, DASD fast write, and cache fast write cache hits.

The number is calculated by dividing the total data set cache hits by the total cache candidate I/O count for the group.

The data set with the highest "Lst%" value is getting the best performance benefits from cache.

Data sets with very large “I/O%” values, but very low cache hit “Lst%” values, could be moved to paths without cache controllers.

For a more detailed analysis of cache activity, use the Cache Manager component.

Column m1

-Cache Hits-
Dsn%

Note: The abbreviation in this column heading matches the value on the screen’s “Detail” field. In this example, the abbreviation “Dsn%” corresponds to the detail value of “Dataset.”

This column shows what percentage of each data set’s total cache candidate I/O count consists of read, read sequential, DASD fast write, and cache fast write cache hits.

The number is calculated by dividing each data set’s total cache hit count by its total cache candidate I/O count.

Data sets that have low cache hit “Dsn%” values are not utilizing cache efficiently.

Column i8

--I/O Time Benefit---
Saved

This column shows the total I/O time-savings (in milliseconds, seconds, minutes, or hours) that resulted from cache hits.

This number is calculated by multiplying the total number of hits by the average disconnect time for each miss.

A large number of hits indicate that the benefits of using cache are high, because DASD accesses are eliminated and I/O time is saved.

Notes:

- If you are not authorized to use the Cache Manager Component of CA-ASTEX, “N/A” appears in this column.
- This number (as well as columns j8 and k8) displays “RLC” for controllers that support record-level caching since CA-ASTEX cannot accurately calculate the cost of using cache.

Column j8

--I/O Time Benefit---
Cost

This column shows the cost of using cache. The cost of using cache is the extra device and path busy time that results from loading the balance of a track after a cache miss.

This number is the total additional connect time needed to load the balance of the track for each cache miss.

A high cost is justified only if using cache also results in a large time-savings ("Column i8") and net benefit ("Column k8").

Notes:

- For read sequential I/Os, no additional cost is incurred when the BLKSIZE-BUFNO combination is larger than 2 tracks for a 3880 cache controller or larger than 3 tracks for a 3990 cache controller.
- If you are not authorized to use the Cache Manager Component of CA-ASTEX, "N/A" appears in this column.
- This number (as well as columns i8 and k8) displays "RLC" for controllers that support record-level caching since CA-ASTEX cannot accurately calculate the cost of using cache.

Column k8

--I/O Time Benefit---
Net

This column shows the net savings from using cache.

This number is calculated by subtracting "I/O time benefit cost" from "I/O time benefit saved."

This is the best measure of overall cache benefit. The larger the net benefit, the more effectively cache is being utilized. If the cost of using cache becomes greater than the savings, net benefit is negative. Volumes and data sets with negative net benefit adversely affect the cache subsystem; you should remove these volumes and data sets from cache.

Notes:

- To remove a data set from cache, you move it to a non-cached volume or controller. If you are using SMS, you can turn off cache for the data set instead of moving it.
- If you are not authorized to use the Cache Manager Component of CA-ASTEX, "N/A" appears in this column.

- This number (as well as columns i8 and j8) displays “RLC” for controllers that support record-level caching since CA-ASTEX cannot accurately calculate the cost of using cache.

Interpreting SMS Problem Analysis Data

What Is the SMS Problem Analysis Screen?

The SMS Problem Analysis screens and reports show the I/O performance of data sets and millisecond response objectives for the must-cache, may-cache, and never-cache storage classes. Using the statistics on these screens and reports, you can compare a data set's actual performance against the must-cache, may-cache, and never-cache objectives that are set for it in SMS.

This chapter describes the columns of data found on the "SMS Problem Analysis" screens. The following information is provided (where applicable) for each column:

- Description of the column or group of columns
- How DASD Manager calculates the value
- How you can use the information

A sample SMS Prob Analysis screen is shown below. As you read the description of each column heading, refer to this example.

```

Summary: System View: Contention      Interval 9:00 - 10:00 as of 9:56
Detail: Group Screen: SMS Prob Analysis System: IPOX Date: 24 JAN 01
>
-
Row 1 to 5 of 5
Summary: System View: Contention      Interval 9:00 - 10:00 as of 9:56
Detail: Group Screen: SMS Prob Analysis System: IPOX Date: 24 JAN 01
>
-
I/O ---Prob%--- -----MSR Objectives----- -Class Problem Dist%-
Group % Sys% Grp% Must May Never NSMS OBJ% Must May Never NSMS
.a.....c.....e.....f.....g2...h2...i2....n....o....a3...b3...c3....d3.
System 100.0 21.2 20.6 12 26 35 0 90 68 14 2 16
PATH3990 85.7 17.6 20.6 12 26 35 0 90 82 16 2 0
PATH02 8.0 2.5 31.3 0 0 0 30 90 0 0 0 100
PATH01 4.8 0.5 9.7 12 26 35 30 90 1 0 0 99
PATH04 1.5 0.7 43.2 0 0 0 30 90 0 0 0 100
PATH03 0.0 0.0 0.0 12 26 35 30 90 0 0 0 0
***** BOTTOM OF DATA *****

```

Understanding the Column Headings

Following are description of the column headings that appear on the “SMS Problem Analysis” screen.

Columns x1 and y1

(Volser) (Jobname)

These columns are the same as the “x1” and “y1” columns that appear on the Problem Analysis Screen.

Column a

summary/detail value

This column identifies the specific “summary” and “detail” data being displayed. In the sample “[SMS Prob Analysis screen](#),” the “summary” data shows system-wide measurements, while the “detail” data shows the performance of contention groups: PATH3990, PATH02, and so on.

The first horizontal row of data conveys information at the “Summary” level. The remaining rows of the screen contain “Detail” level data. If the “Summary” value were Group, it would list the names of volumes belonging to a particular storage or contention group.

An asterisk (*) after each volser means that CA-ASTEX is currently collecting detailed job and data set data for this volume. A plus sign (+) indicates that exception mode is specified and that job and data set data was collected at some point during the current interval but is not being collected now.

If the “Detail” value on this screen was System, and you were using a merged Interval Database as input, this screen would list the systems for your complex.

Column b

Total
I/O

This column shows the total number of I/O requests issued within the group or volume and for each data set or job.

Note: This heading appears only on screens with a detail field value of Job or Dataset.

Columns c, e, and f

I/O --Prob%--
% Sys% Grp%

These columns are the same as the “c,” “e,” and “f” columns that appear on the Problem Analysis screen.

Note: These fields only appear on certain screens.

Column s2

-----Must%-----
I/O

This column shows the percentage of total I/O requests that were directed to must-cache data sets.

The number is calculated by dividing the number of must-cache I/Os by the total number of I/Os.

Note: This column appears only on an SMS Problem Analysis screen with a “Detail” value of Dataset, Job, or Class.

Column j2

-----Must%-----
Prb

This column shows the percentage of must-cache I/Os whose response times did not meet your defined MUSTMSR objective.

This number is calculated by dividing the number of must-cache response problems by the total number of must-cache I/Os.

A high value in this column indicates that must-cache data set I/Os are being delayed. Use the “Problem Analysis” screen to determine what types of problems you are having and resolve them.

Note: This column appears only on an SMS Problem Analysis screen with a “Detail” value of Dataset, Job, or Class.

Column t2

-----Must%-----
Cnd

This column shows the percentage of must-cache I/Os that were cache candidates.

The number is calculated by dividing the number of must-cache I/Os that were cache candidates by the total number of must-cache I/Os.

Note: This column appears only on an SMS Problem Analysis screen with a “Detail” value of Dataset, Job, or Class.

Column u2

-----Must%-----
Hit

This column shows the percentage of must-cache cache candidates that were cache hits.

This number is calculated by dividing the number of must-cache I/O cache hits by the total number of must-cache cache I/O candidates.

A low value in this column indicates that the data set may not be using cache efficiently. Use the Cache Manager component of CA-ASTEX to manage your cache resources.

Note: This column appears only on an SMS Problem Analysis screen with a “Detail” value of Dataset, Job, or Class.

Column v2

-----May%-----
I/O

This column shows the percentage of the total I/O requests that were directed to may-cache data sets.

This number is calculated by dividing the number of may-cache I/Os by the total number of I/Os.

Note: This column appears only on an SMS Problem Analysis screen with a “Detail” value of Dataset, Job, or Class.

Column k2

-----May%-----
Prb

This column shows the percentage of may-cache I/Os whose response times did not meet your defined MAYMSR objective.

The number is calculated by dividing the number of may-cache response problems by the total number of may-cache I/Os.

A high value in this column indicates that may-cache data set I/Os are being delayed. Use the “Problem Analysis” screen to determine what types of problems you are having and resolve them.

Note: This column appears only on an SMS Problem Analysis screen with a “Detail” value of Dataset, Job, or Class.

Column w2

-----May%-----
Cnd

This column shows the percentage of may-cache I/Os that were cache candidates.

The number is calculated by dividing the number of may-cache I/Os that were cache candidates by the total number of may-cache I/Os.

Note: This column appears only on an SMS Problem Analysis screen with a “Detail” value of Dataset, Job, or Class.

Column x2

-----May%-----
Hit

This column shows the percentage of may-cache cache candidates that were cache hits.

The number is calculated by dividing the number of may-cache cache hits by the total number of may-cache cache candidates.

A low value in this column indicates that the data set may not be using cache efficiently. Use the Cache Manager component of CA-ASTEX to manage your cache resources.

Note: This column appears only on an SMS Problem Analysis screen with a “Detail” value of Dataset, Job, or Class.

Column y2

-Never%-
I/O

This column shows the percentage of the total I/O requests that were directed to never-cache data sets.

The number is calculated by dividing the number of never-cache I/Os by the total number of I/Os.

Note: This column appears only on an SMS Problem Analysis screen with a “Detail” value of Dataset, Job, or Class.

Column l2

-Never%-
Prb

This column shows the percentage of never-cache I/Os whose response times did not meet your defined NEVMSR objective.

The number is calculated by dividing the number of never-cache response problems by the total number of never-cache I/Os.

A high value in this column indicates that never-cache data set I/O's are being delayed. Use the "Problem Analysis" screen to determine what types of problems you are having and resolve them.

Note: This column appears only on an SMS Problem Analysis screen with a "Detail" value of Dataset, Job, or Class.

Column g2

-----MSR Objectives-----
Must

This column shows the millisecond response objective for must-cache data sets. This percentage comes from the MUSTMSR parameter in your CA-ASTEX parameter data set.

Note: This column appears only on group or volume screens.

Column h2

-----MSR Objectives-----
May

This column shows the millisecond response objective for may-cache data sets. This percentage comes from the MAYMSR parameter in your CA-ASTEX parameter data set.

Note: This column appears only on group or volume screens.

Column i2

-----MSR Objectives-----
Never

This column shows the millisecond response objective for never-cache data sets. This percentage comes from the NEVMSR parameter in your CA-ASTEX parameter data set.

Note: This column appears only on group or volume screens.

Column n

-----MSR Objectives-----
NSMS

This column shows the millisecond response objective for non-SMS data sets. This percentage comes from the NSMSMSR parameter in your CA-ASTEX parameter data set.

Note: This column appears only on group or volume screens.

Column o

-----MSR Objectives-----
OBJ%

This column displays the percentage of time your response time objectives should be met. This percentage comes from the MSROBJ parameter in your CA-ASTEX parameter data set.

Note: This column appears only on group or volume screens.

Column a3

-Class Problem Dist%-
Must

This column shows the percentage of total response exceptions that were due to must-cache I/Os missing their MSR objective.

This number is calculated by dividing the number of must-cache response problems by the total number of response problems.

A high value in this column indicates that must-cache data set I/O's are being delayed. Use the Problem Analysis screen to determine what types of problems you are having and resolve them.

Column b3

-Class Problem Dist%-
May

This column shows the percentage of total response exceptions that were due to may-cache I/Os missing their MSR objective.

This number is calculated by dividing the number of may-cache response problems by the total number of response problems.

A high value in this column indicates that may-cache data set I/O's are being delayed. Use the Problem Analysis screen to determine what types of problems you are having and resolve them.

Column c3

-Class Problem Dist%-
Never

This column shows the percentage of total response exceptions that were due to never-cache I/Os missing their MSR objective.

This number is calculated by dividing the total number of never-cache response problems by the total number of response problems.

A high value in this column indicates that never-cache data set I/O's are being delayed. Use the Problem Analysis screen to determine what types of problems you are having and resolve them.

Column d3

-Class Problem Dist%-
NSMS

This column shows the percentage of total response exceptions that were due to non-SMS I/Os missing their MSR objective.

This number is calculated by dividing the number of non-SMS response problems by the total number of response problems.

A high value in this column indicates that non-SMS data set I/Os are being delayed. Use the Problem Analysis screen to determine what types of problems you are having and resolve them.

Interpreting Reserve-Allocation Data

Understanding the Column Headings

The Reserve-Allocation screens and reports show reserve performance measurements for individual volumes and groups of volumes. These screens and reports also show the Allocation Director activity for these volumes and groups of volumes. You can use these screens and reports to determine why particular volumes or groups may be experiencing response time delays due to reserve problems and review how the Allocation Director has helped.

The “Reserve-Allocation” screens and reports repeat certain I/O measurements from the “Problem” and “Response Analysis” screens while providing additional reserve and allocation data.

This chapter describes the columns of data found on the “Reserve-Allocation” screens and reports. The following information is provided (where applicable) for each column:

- Description of the column or group of columns
- How DASD Manager calculates the value
- How you can use the information

A sample Reserve-Allocation screen is shown below. As you read the description of each column heading, refer to this example.

```

Row 1 to 10 of 10
Summary: Group   View : Contention   Interval 15:06 - 15:30 as of 15:07
Detail : Volume  Screen: Reserve-Allocation System: S028 Date: 26 JAN 01
>
Group      I/O  Rsv Non-Busy Alloc -Alloc Dist- --Allocation Distribution%--
Group      %   %   %   Total Sum% Rcll% Dir To Dir From SRM Unbal
. .a.....c.....d1....e1.....m2.....o2...o3.....p2.....q2.....r2....a1.
TSUGRP 100.0 0.7 2.6 8,000 100.0 40.3 25.0 28.0 13.0 34.0
TSU001 64.9 1.0 6.6 2,300 28.8 26.3 68.0 4.0 9.0 9.0
TSU002 30.1 0.7 0.1 3,500 43.8 67.2 17.2 73.0 2.0 8.0
TSU003 5.0 0.3 0.9 2,200 27.5 12.0 3.0 2.0 5.0 90.0
*****BOTTOM OF DATA*****

```

Column a

summary/detail values

This column lists the names (volser) of individual volumes in a particular contention or storage group.

If the “Summary” value were System, it would list the names of groups on the system.

An asterisk (*) after each volser means that CA-ASTEX is currently collecting detailed job and data set data for this volume. A plus sign (+) indicates that exception mode is specified and that job and data set data was collected at some point during the current interval but is not being collected now.

No “Reserve-Allocation” screen exists for data set or job data. However, if you see an “R” in front of a data line on a screen that contains detailed data set data, it means that all I/O requests for that data set occurred while the I/O device was reserved.

Column c

I/O
%

This column is the same as the “c” column from the “Problem Analysis” screen.

Column d1

Rsv
%

This column shows the percentage of time a volume is reserved.

The number is calculated by dividing the volume’s total reserve time by the total time in the specified interval.

The volume that has the highest percentage is the one most often reserved. See the upcoming explanation of the “Non-Busy%” column.

Column e1

Non-Busy
%

“Non-busy” reserve time represents the amount of time an I/O device is reserved but is not actually busy performing an I/O operation. Other systems cannot access the device, and response time is affected.

This is the percentage of time, on average, that the volume was in a non-busy reserve state during the interval. The sample “[Reserve-Allocation screen](#)” shows that volume TSU001 was in a non-busy reserved state 6.6% of the time so far this interval. Overall, volumes in the group spent 2.6% of their time (the number at the top of the column) in non-busy reserved states.

The number is calculated by dividing the sum of all non-busy reserve times for a particular volume by the total amount of time in the specified interval.

You may want to use the CA-ASTEX trace facility to further investigate volumes that have the highest non-busy reserve percentages.

Having a large number of non-busy reserves can seriously affect response time on other CPUs in your complex. If you want to eliminate hardware reserves for data sets that have high non-busy percentages, consider converting those hardware reserves to global ENQUEUEs. (Computer Associates offers a software product, CA-MIM, which has a component called Multi-image Integrity that converts reserves to global ENQUEUEs.)

Column m2

Alloc
Total

This column shows the number of allocation requests made to this volume. This includes the number of allocations for which this volume was selected and the number of allocation requests that were directed away from this volume to another volume.

A high number here indicates that the volume experienced a large amount of activity. Review the "Allocation Distribution%" columns to determine what kind of allocation activity occurred.

Column o2

-Alloc Dist-
Sum%

This column shows the percentage of the group's allocation requests that were contributed by this volume. In the sample "[Reserve-Allocation screen](#)," volume TSU002 contributed 43.8 percent of the total number of allocation requests to the TSUGRP.

This number is calculated by dividing the number of the volume's allocation requests by the total number of allocation requests for the group.

A high number here indicates that the volume has had a large amount of allocation activity. Review the "Allocation Distribution%" columns to determine what kind of allocation activity occurred.

Column o3

-Alloc Dist-
Rc11%

This column shows the percentage of the total number of allocations that were HSM recalls for this volume.

This number is calculated by dividing the number of HSM recalls by the total number of allocations.

A high number here indicates that the volume has had a large amount of HSM recall activity.

Column p2

```
--Allocation Distribution%--  
Dir To
```

This column shows the percentage of allocations directed to this volume.

The number is calculated by dividing the number of allocations directed to this volume by the total number of allocations.

A high percentage here indicates that the Allocation Director has had high involvement with the volume and has been directing allocations from other volumes to this volume.

Column q2

```
--Allocation Distribution%--  
Dir From
```

This column shows the percentage of allocations CA-ASTEX directed from this volume.

This number is calculated by dividing the number of allocations directed from the volume by the total of all allocation requests to this volume.

A high number here indicates that the Allocation Director determined that this volume is a poor allocation candidate and is directing allocation requests to other volumes.

If CA-ASTEX had not determined that the volume was a poor allocation candidate and had not directed allocations away from the volume, performance problems could have occurred.

Column r2

```
--Allocation Distribution%--  
SRM
```

This column shows the percentage of allocations that CA-ASTEX determined were good allocation decisions from SRM.

This number is calculated by dividing the number of allocations SRM directed to the volume by the total of all allocation requests to this volume.

A high number here indicates that the Allocation Director has determined this volume is a good allocation candidate and has allowed SRM's allocation decisions to be implemented.

Column a1

--Allocation Distribution%--
Unbal

This column shows the percentage of allocations that are going to SMS's secondary volume list. See Finding Out-of-Balance Conditions in SMS in the "Using the Allocation Director" chapter for more detailed information on how SMS handles allocations.

Note: In a non-SMS environment, this column shows how many allocations to the particular group or volume had to be re-attempted.

This number is calculated by dividing the number of allocations made using the secondary list by the total number of allocation requests.

A high percentage in this column indicates that most allocation candidates are in the secondary list. As a result, an unbalanced condition may arise, causing your resources to be poorly utilized.

To help balance the allocations, you may want to examine your storage class MSR values and your free space thresholds under ISMF. To alleviate space constraints, consider using the HSM interval migration.

Interpreting Seek Analysis Data

How to Access the Seek Analysis Screen

The Seek Analysis screen is accessed differently than the Storage Statistics screens. You *must* be in the Reserve Allocation screen to access the Seek Analysis screen. From the Reserve Allocation screen, you can:

- Enter a D in the Selection column next to the desired Summary value to see a Detail Dataset Seek Analysis screen.
- Enter a J in the Selection column next to the desired Summary value to see a Detail Jobname Seek Analysis screen.

Understanding the Column Headings

The Seek Analysis screens and reports show statistics for either data sets or jobs. You can use these screens and reports to identify the data sets or jobs whose high seek rates are contributing to poor response time on a volume.

This section describes the columns of data found on the “Seek Analysis” screens and reports. The following information is provided (where applicable) for each column:

- Description of the column or group of columns
- How DASD Manager calculates the value
- How you can use the information

A sample Seek Analysis screen is shown below. As you read the description of each column heading, refer to this example.

Summary: Volume View: Contention Interval 9:00 - 10:00 as of 9:56											
Detail: Dataset Screen: Seek Analysis System: IPOX Date: 24 JAN 01											
> Scroll HALF											
-											
Total I/O ---Total Seek--- ---Inter Seek--- ---Intra Seek---											
Dataset	I/O	%	Avg	Lst%	Dsn%	Avg	Lst%	Dsn%	Avg	Lst%	Dsn%
..y1..a....b....c....o1....p1....q1....r1....s1....t1....u1....v1....w1.											
SYS808	39,486	100.0	43	26.3		320	2.9		8	23.4	
X EIMSVS.RESLIB											
RC10DL	19,871	50.3	12	21.5	42.7	220	0.4	0.8	9	21.1	42.0
EIMSVS.TFORMAT											
DA18D47	11,040	28.0	93	1.7	6.0	283	0.5	1.8	10	1.2	4.2
EIMSVS.USERRSLB											
DSIDT11	2,611	6.6	226	0.5	7.9	406	0.3	4.3	5	0.2	3.5

Columns x1 and y1

(Volser) (Jobname)

These columns are the same as the “x1” and “y1” columns on the “Problem Analysis” screen.

Column a

summary/detail values

This column contains the names of the most active data sets on a selected volume (SYS808, in this example). Sometimes, the catchall name “***OTHER” appears in this column. CA-ASTEX uses it for I/O statistics that do not pertain to any one data set, for example, a stand-alone release CCW that releases a device reserve.

The name “***TEMP” also may appear in this column. CA-ASTEX uses it as a generic name for all temporary data sets.

On a screen of job data, this column would list the names of the most active jobs on the volume.

Note: An “X” to the left of a data set name indicates that cross system seeks have occurred for this data set. An “X” appears only if 33 percent of all seeks were cross-system seeks.

An “R” to the left of a data set name means that all I/O requests for that data set occurred while the I/O device was reserved.

An “S” to the left of a data set name means that the data set has been defined for sequential data set striping.

An “H” to the left of a data set name means that the data set has been defined as a hierarchical file system (HFS) data set.

An “SR” to the left of a data set name means that all I/O requests for that data set occurred while the I/O device was reserved and that the data set has been defined for sequential data set striping.

An “HR” to the left of a data set name means that all I/O requests for that data set occurred while the I/O device was reserved and that the data set has been defined as a hierarchical file system (HFS) data set.

An “HS” to the left of a data set name means that the data set has been defined for sequential data set striping and that the data set has been defined as a hierarchical file system (HFS) data set.

A “QR” to the left of a data set name means that all I/O requests for that data set occurred while the I/O device was reserved, that the data set has been defined for sequential data set striping, and that the data set has been defined as a hierarchical file system (HFS) data set.

Columns b and c

Total I/O	I/O %
--------------	----------

The “Total I/O” counts are the same as Column “b” on the “Problem Analysis” screen.

The “I/O%” values are the same as Column “c” on the “Problem Analysis” screen.

Column o1

---Total Seek---
Avg

A “seek” is an I/O request that causes a DASD arm to move at least one cylinder’s length.

This column shows the average distance a DASD arm traveled during a series of seeks. (This includes seeks associated with a single data set or job and seeks that occur for an entire group of data sets or jobs.)

The number is calculated by dividing the total number of cylinders traversed by the total number of seeks.

Column p1

---Total Seek---
Lst%

This column shows what percentage of the total I/O count for all data sets listed on the screen consisted of seeks.

The number is calculated by dividing the total number of seeks for each data set by the total I/O count for all data sets listed on the screen.

The data set with the largest LST% value contributes the most to the volume's seek activity. Examine data in the intra- and inter-seek columns to further identify the seek problem.

Column q1

---Total Seek---
Dsn%

Note: If the detail field value is Job, the abbreviation in this column heading is "Job%" instead of "Dsn%."

This column shows what percentage of each data set's total I/O count consisted of seeks (I/Os that required a DASD arm to move at least one cylinder's length).

The number is calculated by dividing the total number of seeks for each data set by its total I/O count.

Data sets that have high total seek percentages have excessive seek activity. Examine data in the intra- and inter-seek columns to further identify the seek problem.

Column r1

---Inter Seek---
Avg

"Inter-seeks" are seeks that occur for different members of a group of data sets or jobs. Suppose a group of data sets received a total of five I/O requests, and three of those requests were "seeks." If the first seek was for data set "A," the second for data set "B," and the third for data set "C," then all three would be "inter-seeks" because each involved a different data set.

This column shows the average distance the DASD arm traveled while doing seeks for a series of different data sets (or different jobs). As shown in the sample ["Seek Analysis screen,"](#) the average inter-seek distance for I/O requests issued to data set EIMSVS.RESLIB is 220 cylinders. For EIMSVS.TFORMAT, the average is 283 cylinders.

The number is calculated by dividing the total number of cylinders traversed during inter-seeks by the total number of inter-seeks that occurred.

Use the seek analysis routines to reorganize volumes with either FDR/COMPAKTOR or DFDSS.

Column s1

---Inter Seek---
Lst%

This column shows the percentage of inter-seeks each data set had with respect to the Total I/O count (39,486) for all the data sets listed. The sample "[Seek Analysis screen](#)" shows that 2.9% of all I/Os for data sets on this volume were inter-seeks. 0.5% of these inter-seeks were associated with the data set EIMSVS.TFORMAT.

The number is calculated by dividing the number of inter-seeks for each data set by the total number of I/O requests for all the data sets listed.

Investigate the data set with the highest inter-seek average. It may be poorly placed, and you may need to reorganize the volume. Use the CA-ASTEX trace facility to gather further information about such data sets.

Column t1

---Inter Seek---
Dsn%

Note: If the detail field value is Job, the abbreviation in this column heading is "Job%" instead of "Dsn%."

This column shows what percentage of each data set's total I/O requests were inter-seeks. In the sample "[Seek Analysis screen](#)," 4.3% of the EIMSVS.USERRSLB data set's total I/O requests were inter-seeks.

The number is calculated by dividing the number of inter-seeks this data set had by its Total I/O count.

Investigate the data set with the highest inter-seek average. That data set may be poorly placed, and you may need to reorganize the volume.

Column u1

---Intra Seek---
Avg

"Intra-seeks" are a series of seeks associated with a single data set or job. Suppose that a data set named "SYS1.CMNDLIB" receives five I/O requests, and no other I/O requests are issued to the volume where it resides. If the first, second, and fourth I/O requests each require DASD arm movement of at least one cylinder, then all three are "intra-seek" requests.

This column shows the average distance the DASD arm traveled during a series of intra-seek requests. The sample "[Seek Analysis screen](#)" shows that the average intra-seek distance for data sets on volume SYS808 is 8 (the number at the top of the column). The average intra-seek distance on data set EIMSVS.TFORMAT is 10 cylinder lengths.

The number is calculated by dividing the total number of cylinders traversed during a series of intra-seeks by the total number of intra-seeks that occurred.

Data sets with high intra-seek averages may need to be reorganized. Use the CA-ASTEX trace facility to further investigate such data sets.

Column v1

---Intra Seek---
Lst%

This column shows the percentage of intra-seek requests each data set had with respect to the total I/O count (39,486) for all the data sets listed. The sample [“Seek Analysis screen”](#) shows that 23.4% of all I/O requests to data sets in this screen were intra-seeks. The data set EIMSVS.RESLIB was associated with 21.1% of these intra-seeks.

The number is calculated by dividing the number of intra-seeks for each data set by the total number of I/O requests for all the data sets included in this screen (39,486).

Investigate the data set with the highest intra-seek average. It may need to be reorganized. Use the CA-ASTEX trace facility to gather further information about such data sets.

Column w1

---Intra Seek---
Dsn%

Note: If the detail field value is Job, the abbreviation in this column heading is “Job%” instead of “Dsn%.”

This column shows what percentage of I/O requests issued to each data set was intra-seeks. In the sample [“Seek Analysis screen”](#) 42.0% of the I/O requests to the EIMSVS.RESLIB data set were intra-seeks.

The number is calculated by dividing the number of intra-seeks this data set had by its Total I/O count.

Investigate the data set with the highest intra-seek average. You may need to reorganize this data set. A large “Dsn%” value may result when a data set is in multiple extents.

Displaying the System Summary Screen

What Is the System Summary Screen?

If you want to see average performance measurements for the entire system before deciding what data to display first, you can look at the “System Summary” screen. This screen, shown below, provides an overview of system-wide DASD performance measurements collected over the indicated period of time (either “All-day” or “Interval”).

Note: This screen is available only as an online screen in the Online Storage Statistics facility. If you are using a merged Interval Database as input, this screen has a summary level of “Complex” and the data displayed applies to the complex.

CA-ASTEX 2.8			
Summary: System		Interval: 9:00 - 10:00 as of 9:56	
Screen: System Summary		System: IPOX	Date: 24 JAN 01
Total I/Os	: 893,140	Per Second	: 265.8
Rsp Avg	: 17	ESCON DPB	: 1.1
Rsp IOSQ	: 1.3		
Rsp Connect	: 6.4	Rsp Pending	: 0.2
Rsp Disconnect	: 8.3	Rsp Dispatch	: 0.7
Prob%	: 10.6	Prob - PDC%	: 0
Prob - Path%	: 23	Prob - Xsy%	: 3
(RPS%)	: 36	Prob - Vol%	: 26
(ESCON%)	: 64	Prob - Seek%	: 8
Prob - Dsp%	: 1	Prob - Dsn%	: 39
Prob - Must%	: 62	Prob - Never%	: 0
Prob - May%	: 33	Prob - NonSMS%	: 5
Cache Cand%	: 72.1	Cache Hits%	: 85.2
RSV Time%	: 3.3	Non-busy RSV Time%	: 2.7
F1= Help F3= End	Press ENTER	to continue	

How to Access the System Summary Screen

This screen is accessed differently than the storage statistics screens you were introduced to in the "Displaying Storage Statistics" chapter. To access this screen, you can:

1. Issue the SYSTEM command from the DASD Manager screen's command area.
2. On the "Select Defaults" screen, specify that this screen is displayed first when you enter the Online Storage Statistics facility. To leave this screen, you either press ENTER or F3/15. If this was the first screen you saw after you left the "Select Input" panel, you advance to a "Problem Analysis" screen. Otherwise, you return to the screen where you issued the SYSTEM command.

Displaying the Data Set Analysis Screen

What Is the Data Set Analysis Screen?

If you want to see detailed I/O statistics for a specific data set that seems to be experiencing poor response, you can look at the “Data Set Analysis” screen shown below. This screen helps you determine whether certain characteristics of a data set or the way it is accessed cause problems, regardless of its placement on a path or volume.

Note: This data is available only as a screen in the Online Storage Statistics facility.

```
Summary: Dataset View : Storage          Interval 9:00 - 10:00 as of 9:56
          Screen: Data Set Analysis      System: IPOX   Date: 24 JAN 01
>
-
Dataset   Total   % Dsn   ---Dsn Problem Distribution---
Volume    I/O     Probs   Intra Seek%  DsnQ%    Other%  Dsorg  Last Job
...a.....b.....a2.....b2.....c2.....d2.....e2.....y1...
ISR.PRE.ISRTLIB
MVSNHJ    3,625    62      15          65       20      P0     PREIN66

Program Management Components:          Intra Seek Components:

% of Total Dsn I/O      =      0      Avg Intra Seek Distance = 256
Bldl/Find %             =      0      Intra Seek % of Dsn I/O = 32.0
Fetch %                  =      0      No. of Allocated Extents = 4

IOS Queue Components:                  Other Components:

Avg IOS Queue Time      = 2.7      Avg Latency+Connect Time = 8.1
                               Avg Latency Time      = 6.7
                               Avg Connect Time       = 1.4
```

How to Access the Data Set Analysis Screen

This screen is accessed differently than the storage statistics screens you were introduced to in the "Displaying Storage Statistics" chapter. To access this screen, you must:

1. Be on a screen where the "Detail" field value is Data set.
2. Enter a D or S in the "Selection column" to the left of one of the data set names.

To leave the "Data Set Analysis" screen, press either ENTER or F3/15, and you return to the screen where you selected it.

Note: You can issue only TSO and ISPF commands from the command area on this screen. The quickest way to update the screen is to return to the previous screen by pressing F3, issue a RESHOW command there, and then select this screen again.

Understanding The Column Headings

Following are descriptions of the column headings that appear on the "Data Set Analysis" screen. The following information is provided (where applicable) for each column:

- Description of the column (or group of columns)
- How DASD Manager calculates the value
- How you can use the information

Column a

Dataset
Volume

The first information that appears beneath this column heading is the name of the data set that you have requested more information about. In the sample "[Data Set Analysis](#)" screen, the data set being highlighted is ISR.PRE.ISRTLIB.

Beneath the data set name, a six-character volser identifies the volume ("MVSNHJ" in this example) where the data set resides.

Column b

Total
I/O

This column shows the total number of I/O requests issued for this data set during the interval of time indicated at the top of the screen.

In the sample “[Data Set Analysis](#)” screen, the ISR.PRE.ISRTLIB data set had 3,625 I/O requests between 9:00 a.m. and 9:56 a.m.

Column a2

% Dsn
Probs

This column shows what percentage of the data set’s total I/O requests were data set problems.

The number is calculated by dividing the number of data set problems for this data set by its total I/O count.

This percentage of data set problems reflects the severity of the data set’s I/O problems. For example, a value of 100% means that every I/O request was a problem. However, a value of 1% means that very few I/O requests became data set problems.

You should make tuning adjustments for data sets that have large data set problem percentages and total I/O counts.

Column b2

---Dsn Problem Distribution---
Intra Seek%

This column shows what percentage of the data set problems resulted from unacceptably long seek times within the data set (intra-seek problems).

The number is calculated by dividing the number of intra-seek problems for this data set by its total number of data set problems.

Intra-seek problems often occur when a data set occupies several non-adjacent extents on a DASD. The random access of large data sets also causes intra-seek problems.

To reduce intra-seek problems, make sure the data set occupies only one extent. Use the Reorganization Utilities to reorganize the contents of both partitioned and non-partitioned data sets. You also may find useful information about data set I/Os in a printout from the CA-ASTEX trace facility.

Column c2

---Dsn Problem Distribution---
DsnQ%

This column shows what percentage of the data set problems resulted because jobs wanting access waited in a queue for another I/O operation on the data set to finish.

The number is calculated by dividing the number of data set queue problems by the total number of data set problems.

Column “y1” shows the name of the last job that accessed this data set. This job’s I/O request may have forced many other jobs to wait for access.

To solve data set queue problems, you must reduce the number of attempts to access the data set. You can divide the data set into several smaller data sets.

Column d2

---Dsn Problem Distribution---
Other%

This column shows what percentage of the data set problems were not attributed to intra-seek or data set queuing problems.

The number is calculated by dividing the remaining number of data set problems (those not attributed to intra-seek or data set queuing problems) by the total number of data set problems.

Most of these “other” types of data set problems occur when the sum of latency and connect times is high. A high latency time may indicate an inefficiently written channel program. A high connect time may result from large block sizes or from excessive searches (such as PDS directory searches or VTOC searches).

Column e2

Dsorg

This column describes the data set's organization.

The value under this heading is one of the following:

CATL	- for catalog data sets
DA	- for BDAM data sets
ISAM	- for ISAM data sets
PO	- for partitioned data sets
PS	- for QSAM data sets
VSAM	- for VSAM data sets
VTOC	- for VTOC data sets
****	- for an unidentified organization

Column y1

Last job

This column shows the last job to access this data set.

In the sample "[Data Set Analysis](#)" screen, the last job was PREIN66.

Program Management Components
% of Total Dsn I/O

This field shows what percentage of the data set's I/O requests were issued from BLDL, FIND, and program fetch routines.

The number is calculated by adding the BLDL, FIND, and fetch I/Os and then dividing that sum by the data set's total I/O count.

For partitioned data sets, this value reflects the amount of actual I/O time spent in finding and loading load modules.

Bldl/Find %

This field shows what percentage of the data set's I/O requests involved partitioned data set directory searches.

The number is calculated by adding the BLDL and FIND I/Os and then dividing that sum by the data set's total I/O count.

This value reflects the amount of actual I/O time you could save by using a software cache for BLDL and FIND calls.

Fetch %

This field shows what percentage of the dataset's I/O requests are issued by a program fetch (to load a module from a partitioned data set).

The number is calculated by dividing the data set's program fetch I/O count by its total I/O count.

This value reflects the amount of actual I/O time you could save by using a software cache for fetches. (Computer Associates' CA-Quick-Fetch product is an example of such software.)

Intra Seek Components

Avg Intra Seek Distance

This field shows the average number of cylinders traversed during an intra data set seek (when a DASD arm moves from one cylinder to another within the data set).

The number is calculated by dividing the number of cylinders traversed during intra-seeks by the number of intra-seek operations.

Longer average intra-seek distances result in longer seek time.

To reduce intra-seek distances, make sure the data set occupies only one extent. Use the Reorganization Utilities to reorganize the contents of both partitioned and non-partitioned data sets. You also may find useful information about data set I/Os in a printout from the CA-ASTEX trace facility.

Intra Seek % of Dsn I/O

This field shows what percentage of the data set's I/O requests involved intra data set seeks.

The number is calculated by dividing the number of intra-seek operations the data set had by its total I/O count.

This value reflects the severity of the data set's intra-seeking problems.

To reduce intra-seeking, make sure the data set occupies only one extent. Use the Reorganization Utilities to reorganize the contents of both partitioned and non-partitioned data sets. You also may find useful information about data set I/Os in a printout from the CA-ASTEX trace facility.

No. of Allocated Extents

This field shows the number of extents allocated for the data set.

This number can help you determine whether intra-seek problems are occurring because the data set occupies multiple extents.

IOS Queue Components

Avg IOS Queue time

Average IOS Queue time is the average time (in milliseconds) that elapses from the time a user program issues an I/O request until the channel subsystem accepts the "Start subchannel."

The number is calculated by dividing the sum of all queue times for the data set by the data set's total I/O count.

Data sets that have consistently high queue times have too much I/O traffic.

Use the Reorganization Utilities to reorganize the contents of both partitioned and non-partitioned data sets. You also may find useful information about data set I/Os in a printout from the CA-ASTEX trace facility.

Other Components

Avg Latency+Connect Time

This field shows the average amount of time (in milliseconds) needed to search for the track sector where the data set resides, to search for records, and to transfer data during an I/O operation for this data set.

The number is calculated by adding the data set's connect and disconnect times then subtracting from that the sum of its RPS and seek times.

If this time seems long, check to see which component, latency time or connect time, is contributing the most to this value.

Avg Latency Time

This field shows the average amount of time (in milliseconds) needed to search for the sector of the track where the data set resides.

The number is calculated by dividing the data set's total latency time by its total I/O count.

Normally, this value should almost equal the time needed to rotate the disk one half-revolution. An inefficient channel program or an excessive amount of cross-system seeking could account for a high average latency time.

The field Avg Latency Time is "N/A" for data sets on RAID 5 and RAID 6 devices.

Avg Connect Time

This field shows the average amount of time (in milliseconds) needed to search for records and transfer data during an I/O operation for this data set.

The number is calculated by dividing the data set's total connect time by its total I/O count.

A large block size, chained reads, or an excessive amount of searching could account for a high average connect time.

Displaying the PAV Analysis Screen

What is the PAV Analysis Screen?

The PAV Analysis Screen shows Parallel Access Volume (PAV) exposure counts. An “exposure” is a UCB associated with the PAV that can be used by IOS to drive an I/O. One Base UCB is online and 0-255 associated Alias UCBs are offline. The WLM-management status of the PAV is also displayed. When viewing realtime data, a list of DUAs for associated Alias UCBs is also displayed.

Notes:

The PAV data in the screen below is available only in the Online Storage Statistics facility.

Data columns in the sample PAV Analysis screen shown below are documented in the section [“Understanding the Column Headings”](#) later in this chapter:

Summary: Volume		View : Contention	Interval 15:01 - 15:10 as of 15:09	
>		Screen: PAV Analysis	System: XE44 Date: 25 JUN 01	
		PAV Exposures	I/O Rsp	----Response Components----- ESCON
Volume	DUA	Cur Max Min	SEC Avg IOSQ Pend Conn Disc Dsp DPB	
....a.....m.....pa.....d.....q.....s.....t.....u.....v.....w.....z...				
MVPC44 *	4203	2 2 2	6.7 2 0.0 0.2 1.4 0.0 0.0 0.0	
DUAs of current Alias UCBs:				(w = managed by WLM)
42FC				

How to Access the PAV Analysis Screen

This screen cannot be scrolled to from other screens. To access this screen, you must:

1. Be on a screen where the “Detail” level is Volume.
2. Enter line command U in the “Selection column” to the left of a PAV Volume serial.

If the volume you selected for display is not defined as a PAV, you will receive ISPF error message ASX611I.

Use column pa (#PAV) on Response Analysis screens to determine if a volume is defined as a PAV.

A P after a Volume serial at the “Summary” level in a Response Analysis screen also indicates that the volume is defined as a PAV.

Understanding the Column Headings

The following are descriptions of the column headings that appear on the “PAV Analysis” screen.

Column a

Volume

This column identifies the specific Volume data being displayed.

See the chapter “Interpreting Response Analysis Data” for more information concerning this column heading.

Column m

DUA

This column contains the device unit address (08C7, 165F, and so on) of the volume.

Column d

I/O
SEC

This column shows the average number of I/O requests issued each second to the volume.

See the chapter “Interpreting Problem Analysis Data” for more information concerning this column heading.

Column pa

PAV Exposures
Cur

This column shows the current number of exposures for this PAV. A preceding **w** indicates the PAV is WLM-managed.

PAV Exposures
Max

This column shows the maximum number of exposures for this PAV at any given point in time during the interval being displayed.

PAV Exposures
Min

This column shows the minimum number of exposures for this PAV at any given point in time during the interval being displayed.

Column q

Rsp
Avg

This column shows the average time (in milliseconds) needed for an I/O request to finish processing.

See the chapter “Interpreting Response Analysis Data” for more information concerning this column heading.

Column s

---Response Components---
IOSQ

This column shows the average time (in milliseconds) that elapses from the time a user program issues an I/O request until the channel subsystem accepts the “Start Subchannel.”

See the chapter “Interpreting Response Analysis Data,” for more information concerning this column heading.

Column t

---Response Components---
Pend

This column shows the average amount of time an I/O request waits in the channel subsystem queue before the I/O operation actually starts on the device.

See the chapter “Interpreting Response Analysis Data” for more information concerning this column heading.

Column u

---Response Components---
Conn

This column shows the average amount of time an I/O device spent searching for and actually transferring data.

See the chapter “Interpreting Response Analysis Data” for more information concerning this column heading.

Column v

---Response Components---
Disc

This column shows the average amount of time (in milliseconds) an I/O device was disconnected and not transferring or searching for data.

See the chapter “Interpreting Response Analysis Data” for more information concerning this column heading.

Column w

---Response Components---
Dsp

This column shows the average amount of time (in milliseconds) that an I/O operation waited to be dispatched either by IOS (the I/O subsystem) or the channel subsystem, or for the I/O interrupt to be processed.

See the chapter “Interpreting Response Analysis Data” for more information concerning this column heading.

Column z

ESCON
DPB

This column shows the average amount of time an I/O request waits due to ESCON director port busy. High values here indicate that you may be experiencing ESCON director port busy delays.

See the chapter “Interpreting Response Analysis Data” for more information concerning this column heading.

DUAs of current Alias UCBs: (w = managed by WLM)

The DUA of each Alias UCB associated with this PAV will be displayed when viewing realtime data. A w after a DUA indicates an Alias UCB being managed by Work Load Manager (WLM). WLM will dynamically change the managed PAV associated with each managed Alias UCB based on workload.

If the number of Alias UCBs associated with this PAV exceeds 169, then the heading line will not be displayed.

If the number of Alias UCBs associated with this PAV exceeds 208, then the WLM status of each Alias UCB will not be displayed.

The maximum number of Alias UCBs that can be associated with one Base UCB is 255.

Reorganizing Volumes and Data Sets

Understanding the Reorganization Utilities

To use CA-ASTEX Reorganization utilities select “Reorganization Utilities” (option D3) from the CA-ASTEX Primary Menu. The menu shown below is displayed:

```
CA-ASTEX 2.8- Reorganization Utilities Menu

OPTION ==>

  1 Volume Reorganization   - Reorganize the data sets on a volume
  2 Member Reorganization   - Reorganize the members of a PDS
  3 TTR Report               - TTR Analysis of a data set

F1= Help F3= End
```

CA-ASTEX offers three utilities for performing reorganization operations on volumes and data sets.

Utility	Function
Volume Reorganization	Reorganizes entire volumes
Member Reorganization	Reorganizes members of a partitioned data set
TTR Report	Helps in reorganizing the contents of sequential data sets

Reorganizing Volumes

The Volume Reorganization utility helps you reorganize entire volumes. It works with FDR/COMPAKTOR or with DFDSS (DFDSS stands for IBM's Data Facility Data Set Services.). The utility program reads CA-ASTEX interval data and performs a seek analysis of a specific volume's performance data. It uses the results of that analysis as a basis for making organizational changes that minimize seek distances for data sets on a volume and ultimately improve performance.

If you are using FDR/COMPAKTOR, the utility program generates control cards that tell COMPAKTOR how to reorganize a volume's data sets in a more efficient sequence that reduces seek distance. If you are using DFDSS, the utility program minimizes seek distances by dynamically allocating and positioning a volume's data sets.

To reorganize a volume, follow these steps:

Step	Action
1	<p>Decide what volumes you want to reorganize.</p> <p>The Complete Solution Plan can make this decision for you. Otherwise, you can choose the volumes yourself by reviewing the Interval Solutions via the Storage Performance Expert, or reviewing data using the Storage Statistics facilities. As a rule, volumes that contain several <i>active</i> data sets with large seek distances are the best candidates for reorganization. Volumes that are experiencing seeking problems have a high "Sk%" value on the "Problem Analysis" screen.</p>
2	<p>Make sure you have performance data from these volumes available.</p> <p>If you want to collect fresh data for the utility programs to analyze, specify a value of YES on the IDB parameter, or a value other than NONE on the DSNSMFT parameter (in the ASXPARM parameter data set), and run CA-ASTEX in detail mode.</p>
3	<p>Make sure the volume to be reorganized is online, but not allocated.</p>
4	<p>Select option 1 from the Reorganization Utilities Menu.</p>
5	<p>Fill in the "Volume Reorganization" panel.</p> <p>Press F1 for online help on how to fill in this panel.</p> <p>If you are using FDR/COMPAKTOR, see the following section.</p> <p>If you are using DFDSS, see the section titled "Using DFDSS".</p>

Using FDR/COMPAKTOR

If you want to use FDR/COMPAKTOR to reorganize your volumes, follow these steps:

Step	Action
1	Select the “FDR” option on the “Volume Reorganization” panel.
2	<p>If you want FDR to <i>simulate</i> a reorganization, specify the following and press ENTER:</p> <p>Backup: NO Seek Analysis: YES Restore: YES Just SIMULATE (FDR only): YES</p>
3	Check return codes and output produced by the simulation.
4	<p>If you want FDR to actually <i>reorganize</i> the volume, specify the following and press ENTER:</p> <p>Backup: YES Seek Analysis: YES Restore: YES Just SIMULATE (FDR only): NO</p> <p>This deletes and reallocates the volume’s data sets. If the restore step fails, use FDR to restore the backup copy of the volume you saved.</p>
5 (Optional)	Edit and submit the JCL.
6	<p>Use the Storage Statistics facilities to compare data collected before and after the reorganization to determine how much reorganizing the volume has improved performance.</p> <p>Note: Refer to the table in “Unmovable Data Sets” for a list of data sets that cannot be moved using FDR/COMPAKTOR.</p>

Using DFDSS

If you are using DFDSS to reorganize your volumes, follow these steps:

Step	Action
1	Select the “DFDSS” option on the “Volume Reorganization” panel.
2	<p>If you just want to see the sequence set without using DFDSS to restore the volume, specify the following and press ENTER:</p> <p>Backup: NO Seek Analysis: YES Restore: NO</p> <p>DFDSS does not provide a method for previewing the reorganization in advance.</p>

Step	Action
3	<p>If you want DFDSS to actually perform the reorganization, specify the following and press ENTER:</p> <p>Backup: YES Seek Analysis: YES Restore: YES</p> <p>This calls DFDSS to delete data sets from the volume and then dynamically allocate and position data sets on the indicated volume.</p> <p>Important! Do not answer YES to “Restore” unless you are absolutely certain that you want DFDSS to reorganize the volume.</p> <p>If either the seek analysis or restore step fails, run DFDSS, using the following input cards, to restore the backup copy of the volume saved on tape:</p> <pre>DUMP INDD(DUMPTOC) OUTDD(DDDDUMMY) - PURGE DELETE</pre> <pre>IF LASTCC GT 8 - THEN EOJ ELSE - RESTORE INDD(RESTTAPE) OUTDD(DUMPTOC) - REPLACE CATALOG - DATASET (INCLUDE(**))</pre>
4 (Optional)	Edit and submit the JCL.
5	<p>Reply to the following message:</p> <p>ASX770D - Initialization complete for vvvvvv, Reply Y to delete datasets or CANCEL.</p>
6	Check the return codes and output from the reorganization job.
7	<p>Use the Storage Statistics facilities to compare data collected before and after the reorganization to determine how much reorganizing the volume has improved performance.</p> <p>Refer to the table in “Unmovable Data Sets” for a list of data sets that cannot be moved using DFDSS.</p>

Unmovable Data Sets

You cannot use DFDSS or FDR/COMPAKTOR to position data sets that have been allocated as unmovable with absolute track positioning. The following table lists the data sets that are unmovable by either DFDSS or FDR. Unmovable data sets are marked with an “X.”

Data Set	DFDSS	FDR
ISAM	X	X
Absolute BDAM	X	X
VVDS	X	X
Any SYS1	X	
VSAM	X	X

How the ASXSEEKA Program Works with DFDSS

DASD Manager’s seek analysis program ASXSEEKA works differently with DFDSS than it does with FDR. When you specify “DFDSS” on the “Volume Reorganization” panel, the ASXSEEKA program links to ADRDSSU (DFDSS) so that it can use DFDSS to delete all movable data sets (as defined by DFDSS) from the volume you want reorganized. DFDSS gets its input from fields on the “Volume Reorganization” panel and the input data sets identified on the “Input Data Sets” panel.

Because PURGE and DELETE are specified in the ASXSEEKA program, DFDSS un-catalogs and deletes the data sets as it writes them out to DUMPTAPE (specified on a DD DUMMY statement). As an additional safety feature, message ASX770D appears at the operator’s console asking you to confirm the deletion. ASXSEEKA does not proceed until you reply to the message. If there are any data sets on the pack that should not be deleted, such as LINKLIST data sets, use the EXCLUDE feature of DFDSS to remove them from consideration.

ASXSEEKA uses dynamic allocation to place the data sets in the proper order on the volume. Cylinder or track organization and secondary allocation size are preserved. All datasets allocated by ASXSEEKA, or restored by DFDSS, are cataloged.

ASXSEEKA does not reposition VSAM data sets when working with DFDSS.

Note: If the ASXSEEKA program fails (after DFDSS deletes the data sets), you must use DFDSS to restore the volume. To ensure success, the RESTORE should be preceded by a PURGE and DELETE of movable data sets (the same as in the ASXSEEKA step), followed by a RESTORE by data set with the REPLACE and CATALOG parameters.

Reorganizing Members of a PDS

The Member Reorganization utility helps you reorganize the members of a partitioned data set (PDS) to reduce intra-seek distances. This utility program reads CA-ASTEX trace records and generates control cards that you can use in an IEBCOPY job to reposition a partitioned data set's members.

To reorganize partitioned data set members, follow these steps:

Step	Action
1	Decide what partitioned data sets need to be reorganized. The Complete Solution Plan identifies these data sets for you. Or, you can choose the data sets yourself by reviewing the Interval Solutions via the Storage Performance Expert, or reviewing data using the Storage Statistics facilities. The best candidates for reorganization are active partitioned data sets that have large <i>intra</i> -seek values.
2	Trace the data set you wish to reorganize. Refer to the <i>CA-ASTEX User Guide</i> for information on issuing the TRACE command.
3	Select option 2 from the Reorganization Utilities menu.
4	Fill out the "Reorganizing Members of a PDS" panel. Press F1 for online help for this panel.
5	When prompted for your trace data sets on the "Select Input Data Sets" screen, select the data sets you want from the list. If none are listed, you can enter their names on this screen.
6 (Optional)	Edit and submit the JCL.
7	Check the output.

Reorganizing Non-partitioned Data Sets

The TTR Report utility produces a “TTR activity report” to help you reorganize the contents of non-partitioned (sequential and BDAM) data sets. This utility program identifies a data set’s areas of greatest I/O activity in relation to their TTR coordinates, by their displacement in tracks (“TT”) and records (“R”) from the beginning of the data set. The report shows which data set records (indicated by their TTR coordinates) have the greatest percentages of I/O activity in relation to the total number of I/Os to the data set.

To generate a TTR activity report for a particular non-partitioned data set, you must:

Step	Action
1	Issue a TRACE command to trace the data set. Refer to the <i>CA-ASTEX User Guide</i> for more information on the TRACE command.
2	Select option 3, “TTR Activity Report,” from the Reorganization Utilities menu.
3	Fill in the “TTR Activity Report” panel. Press F1 for online help on how to fill in this panel.
4	When prompted for your trace data sets on the “Select Input Data Sets” screen, select the data sets you want from the list. If none are listed, you can enter their names on this screen.
5 (Optional)	Edit and submit the JCL.
6	Review the output.

Storage Statistics Display Commands

Command Quick Reference

The following table shows the function of each display command you can use in the storage statistics displays:

Function	Command
Redisplays the last command you entered.	?
Changes the interval of displayed data from “Interval” to “All-day.” (Valid only when viewing real-time data.)	ALLDAY
Displays only data for paths that are part of a cache subsystem.	CACHE
Changes the interval of displayed data from “All-day” to “Interval” (or vice versa). (Valid only when viewing real-time data.)	CHGI
Causes the Command Input screen to appear.	CMD
Takes you from DASD Manager to Cache Manager.	CSTATS
Lets you dynamically change defaults for problem identification.	DEF
Displays data located near the bottom of the screen.	DOWN
Displays information about the most active data sets on a particular volume or in a particular group. (Valid only on a job data screen.)	DSN
Takes you from Cache Manager to DASD Manager.	DSTATS
Displays data at a less detailed level.	END
Lets you dynamically change the Excluded Data Set List.	EXCL
Displays help for the data screen you are viewing.	HELP
Changes the interval of displayed data from “All-day” to “Interval.” (Valid only when viewing real-time data.)	INTERVAL
Displays information about the most active jobs on a particular volume or in a particular group. (Valid only on a data set screen.)	JOB

Function	Command
Displays the version and PTF level of the CA-ASTEX panels, monitor, and the PMC and MMC address spaces.	LEVEL
Locates a particular system, group, volume, data set, or job name.	LOC or L
Displays a newer interval of data from the interval selection list.	NEXT [# M]
Displays an older interval of data from the interval selection list.	PREV [# M]
Exits from the data display facility.	QUIT
Updates the screen with the most current real-time performance measurements. Pressing ENTER has the same effect.	RESHOW or ENTER key
Displays the indicated type of data.	SCREEN [name]
Scrolls to the left through the sequence of DASD Manager's data screens.	SCRNL
Scrolls to the right through the sequence of DASD Manager's data screens.	SCRNR
Sorts up to three specified columns of data in ascending or descending order.	SORT c1 [/] c2 [/] c3 [/]
Displays the System Summary screen.	SYSTEM
Displays data located near the top of the screen.	UP
Changes the "view" from "contention" to "Storage" (or vice versa).	VIEW

Index

%

% Dsn Prob column, 13-3
% of Total Dsn I/O value, 13-5

?

? command, A-1

A

accessing other components, 3-4
Active Time column, 4-2, 6-6
ALLDAY command, A-1
Allocation Director, 1-2
 and HSM-recalled data sets, 2-2
 and the System Resources Manager (SRM), 2-2
 benefits of, 2-5
 compatibility issues, 2-6
 how it works, 2-1
 using, 2-6
Allocations % column, 10-3
Allocations Total column, 10-3
ARCRDTEXT exit, 2-2, 2-3
ASXSEEKA program, 15-5
Avg Intra Seek Distance value, 13-6
Avg. Connect Time value, 13-8
Avg. IOSQ Time value, 13-7
Avg. Latency Time value, 13-8
Avg. Latency+Connect Time value, 13-7

B

Batch Storage Statistics, 1-2, 4-1
 for multiple systems, 1-4
 sample selections, 4-5
Bldl/Find % value, 13-5

C

Cache Candidates screen, column headings
 I/O Time Benefit Cost, 8-6
 I/O Time Benefit Net, 8-6
 I/O Time Benefit Saved, 8-5
CACHE command, A-1
Cache Hits Dsn% column, 8-5
Cache Hits Lst% column, 8-4
Cache-RAID Analysis screen, column headings
 I/O%, 8-2, 8-4
 IRC Sys%, 8-3
 PDC Prb% Dsn%, 8-3
 PDC Prb% Lst%, 8-2
 summary/detail values, 8-2, 8-4
 Volser and Jobname, 8-2, 8-3
CHGI command, A-1
Class Problem Dist% May column, 9-7
Class Problem Dist% Must column, 9-7
Class Problem Dist% Never column, 9-8
Class Problem Dist% NSMS column, 9-8
CMD command, A-1
Column †, 14-3
Command Interface, 3-7

Command Interface screen, 3-4

commands, display

- HELP, A-1
- NEXT, A-2
- RESHOW, A-2
- SCREEN, A-2
- SYSTEM, 12-2, A-2

common

- facilities, 3-7
- measurement routines, 1-4

Connect % column, 7-3

Connect Avg. column, 7-2

connect time, 6-5

contention group, 3-2, 4-2

cross-system reports, 1-4

CSTATS command, A-1

D

data

- analyzed by DASD Manager, 1-4
- trace, 15-6

Data Set Analysis screen, 3-4, 13-1

Data Set Analysis screen, column headings

- % Dsn Prob, 13-3
- % of Total Dsn I/O, 13-5
- Avg. Connect Time, 13-8
- Avg. Intra Seek Distance, 13-6
- Avg. IOSQ Time, 13-7
- Avg. Latency Time, 13-8
- Avg. Latency+Connect Time, 13-7
- Bldl/Find %, 13-5
- Dsn Problem Distribution DsnQ%, 13-4
- Dsn Problem Distribution Intra Seek%, 13-3
- Dsn Problem Distribution Other%, 13-4
- Dsorg, 13-5
- Fetch %, 13-6
- Intra Seek % of Dsn I/O, 13-7
- Last job, 13-5
- No. of Allocated Extents, 13-7
- Total I/O, 13-3

data, trace, 15-7

databases, CA-ASTEX

- trend database, 1-5

DEF command, A-1

DFDSS, 15-2, 15-3

Dir Frm% column, 10-4

Dir To% column, 10-4

disconnect time, 6-5

dispatch time, 6-5, 14-4

DOWN command, A-1

DSN command, A-1

Dsn Problem Distribution DsnQ% column, 13-4

Dsn Problem Distribution Intra Seek% column, 13-3

Dsn Problem Distribution Other% column, 13-4

Dsorg column, 13-5

DSTATS command, A-1

DUA column, 6-2, 14-2

E

END command, A-1

ESCON DPB column, 6-7, 7-4, 14-4

exception, 4-2

EXCL command, A-1

F

FDR/COMPAKTOR, 15-2, 15-3

Fetch % value, 13-6

File Merge/Conversion utility, 1-4, 4-2

function key settings, 3-4

H

HELP command, A-1

How to Access the PAV Analysis Screen, 14-1

HSM and CA-ASTEX's Allocation Director, 2-2

I

I/O % column, 5-3, 7-2, 10-2, 11-3
I/O SEC column, 5-3, 14-2
I/O Time Benefit Cost column, 8-6
I/O Time Benefit Net column, 8-6
I/O Time Benefit Saved column, 8-5
I/O%, 8-2, 8-4
I/O% column, 9-2
IEBCOPY job, 15-6
Inter Seek Avg column, 11-4
Inter Seek Dsn% column, 11-5
Inter Seek Lst% column, 11-5
INTERVAL command, A-1
interval data, 1-4, 4-2
Interval Database, 4-2
Intra Seek % of Dsn I/O value, 13-7
Intra Seek Avg column, 11-5
Intra Seek Dsn% column, 11-6
Intra Seek Lst% column, 11-6
IRC Sys% column, 8-3

J

JCL, 4-3, 15-3, 15-6, 15-7
JOB command, A-1
jobs, IEBCOPY, 15-6

L

Last Access column, 4-2, 6-6
Last job column, 13-5
LEVEL command, A-2
LOC or L command, A-2

M

May% CND column, 9-5
May% Hit column, 9-5
May% I/O column, 9-4
May% Prb column, 9-4
member reorganization utility, 15-6
Mode column, 6-3
moving between screens, 3-4
MSR Objectives May column, 9-6
MSR Objectives Must column, 9-6
MSR Objectives Never column, 9-6
MSR Objectives NSMS column, 9-7
MSR Objectives OBJ% column, 9-7
multiple systems, 1-4
Must% Cnd column, 9-3
Must% Hit column, 9-4
Must% I/O column, 9-3
Must% Prb column, 9-3

N

Never% I/O column, 9-5
Never% Prb column, 9-6
NEXT command, A-2
No. of Allocated Extents value, 13-7
Non-Busy % column, 10-2
non-partitioned data sets
 reorganizing, 15-7

O

Online Cache Statistics, 3-6
Online Storage Statistics, 3-1
 accessing Online Cache Statistics from, 3-7

accessing the Storage Performance Expert from, 3-6
description of, 1-2, 3-1, 4-1
procedure for displaying, 3-3
types of screens, 3-1
online vs batch storage statistics, 4-2

P

Parallel Access Volume (PAV), 14-1

Path-Cache Analysis screen, 3-4

Path-Cache Analysis screen, column headings

- Cache Hits Dsn%, 8-5
- Cache Hits Lst%, 8-4
- Connect %, 7-3
- Connect Avg., 7-2
- ESCON DPB, 7-4
- I/O %, 7-2
- Prb Dist ESC%, 7-5
- Prb Dist RPS%, 7-5
- Pth Prob%, 7-4
- Pth Prob% Dsn%, 7-5
- RPS Avg, 6-6, 7-3
- summary/detail values, 7-2
- Volser and Jobname, 6-2, 7-2

PAV

- Column, 14-2
- Column d, 14-2
- Column m, 14-2
- Column pa, 14-2
- Column q, 14-3
- Column s, 14-3
- Column u, 14-3
- Column v, 14-4
- Column w, 14-4
- Column z, 14-4
- exposure count, 14-1

Prb Dist ESC% column, 7-5

Prb Dist RPS% column, 7-5

PREV [# | M] command, A-2

printing a screen, 3-4

Prob% Grp% column, 5-5, 9-2

Prob% Sys% column, 5-4, 9-2

Problem Analysis screen, 3-4

Problem Analysis screen, column headings

- I/O %, 5-3
- I/O SEC, 5-3, 14-2
- Prob% Grp%, 5-5
- Prob% Sys%, 5-4
- Problem Distribution Pth%, 5-5
- Problem Distribution Vol%, 5-7
- Problem Distribution% Dsn, 5-8
- Problem Distribution% Dsp, 5-7
- Problem Distribution% IRC, 5-6
- Problem Distribution% PDC, 5-6
- Problem Distribution% Sk, 5-7
- Problem Distribution% Xsy, 5-6
- Problem Type, 5-4
- summary/detail values, 5-2
- Total I/O, 5-2

problem discussions

- accessing from screens, 3-6

Problem Distribution% Dsn column, 5-8

Problem Distribution% Dsp column, 5-7

Problem Distribution% IRC column, 5-6

Problem Distribution% PDC column, 5-6

Problem Distribution% Pth column, 5-5

Problem Distribution% Sk column, 5-7

Problem Distribution% Vol column, 5-7

Problem Distribution% Xsy column, 5-6

problem list, accessing from screens, 3-6

Problem Type, 5-4

Program Name column, 6-7

Pth Prob% column, 7-4

Pth Prob% Dsn% column, 7-5

Q

queue time, 6-4, 13-7

QUIT command, A-2

R

R (on data set screens), 10-2, 11-2

real-time data, 4-2

reorganization utilities, 15-1

Reorganization Utilities, 1-2

reorganizing

non-partitioned data sets, 15-7

partitioned data sets, 15-6

volumes, 15-2

Reserve-Allocation screen, 3-4, 10-1

relation to the Allocation DIrector, 2-4

Reserve-Allocation screen, column headings

Allocations %, 10-3

Allocations Total, 10-3

Dir Frm%, 10-4

Dir To%, 10-4

I/O %, 10-2

Non-Busy %, 10-2

Rsv %, 10-2

SRM Accpt%, 10-4

summary/detail values, 10-2

Unbal%, 10-5

Response Analysis screen, 3-4

Response Analysis screen, column headings, 6-7

Active Time, 6-6

DUA, 6-2, 14-2

ESCON DPB, 6-7

Last Access, 6-6

Mode, 6-3

Rsp Avg., 6-3, 14-3

Rsp Components Conn, 6-5

Rsp Components Disc, 6-5

Rsp Components Dsp, 6-5

Rsp Components IOSQ, 6-4, 14-3

Rsp Components Pend, 6-4

summary/detail values, 6-2, 14-2

Response and Exception Analysis report, 4-2

response time

problems, 5-5

response time problems, 5-4

RPS Avg column, 6-6, 7-3

Rsp Avg column, 6-3, 14-3

Rsp Components Conn column, 6-5

Rsp Components Disc column, 6-5

Rsp Components Dsp column, 6-5

Rsp Components IOSQ column, 6-4, 14-3

Rsp Components Pend column, 6-4

Rsv % column, 10-2

S

S (on data set screens), 11-2

SCREEN command, A-2

screen format, 3-4

screens vs reports, 4-2

SCRNL command, A-2

SCRNR command, A-2

seek analysis routine, 15-2

Seek Analysis screen, 3-4

accessing, 11-1

Seek Analysis screen, column headings

I/O %, 11-3

Inter Seek Avg, 11-4

Inter Seek Dsn%, 11-5

Inter Seek Lst%, 11-5

Intra Seek Avg, 11-5

Intra Seek Dsn%, 11-6

Intra Seek Lst%, 11-6

Seek Dsn%, 11-4

Seek Lst%, 11-3

summary/detail values, 11-2

Total I/O, 11-3

Total Seek Avg, 11-3

Volser and Jobname, 11-2

Seek Dsn% column, 11-4

Seek Lst% column, 11-3

selecting a level of detail, 3-4

selecting display defaults, 3-4

selecting expert analysis defaults, 3-4

selecting input data, 3-4

selecting input data sets, 3-4

SMF data, 4-2

SMS Problem Analysis screen, 3-4

SMS Problem Analysis screen, column headings

Class Problem Dist% May, 9-7

Class Problem Dist% Must, 9-7

Class Problem Dist% Never, 9-8

Class Problem Dist% NSMS, 9-8

I/O%, 9-2

May% CND, 9-5

May% Hit, 9-5

May% I/O, 9-4

- May% Prb, 9-4
- MSR Objectives May, 9-6
- MSR Objectives Must, 9-6
- MSR Objectives Never, 9-6
- MSR Objectives NSMS, 9-7
- MSR Objectives OBJ%, 9-7
- Must% Cnd, 9-3
- Must% Hit, 9-4
- Must% I/O, 9-3
- Must% Prb, 9-3
- Never% I/O, 9-5
- Never% Prb, 9-6
- Prob% Grp%, 9-2
- Prob% Sys%, 9-2
- summary/detail values, 9-2
- Total I/O, 9-2
- Volser and Jobname, 9-2
- solution discussions
 - accessing from screens, 3-6
- SORT c1 [/] c2 [/] c3 [/] command, A-2
- SRM Acpt% column, 10-4
- storage group, 3-2, 4-2
- storage hierarchy, 1-1
- Storage Performance Expert
 - accessing from storage statistics screens, 3-6
- summary, system I/O statistics, 12-1
- summary/detail value column, 9-2
- summary/detail values column, 8-2, 8-4, 10-2, 11-2
- switching between views, 3-4
- SYSTEM command, 12-2, A-2
- System Resources Manager
 - and CA-ASTEX's Allocation Director, 2-2
- System Summary screen, 3-4, 12-1

T

- technical support, contacting, 1-5
- Total I/O column, 5-2, 9-2, 11-3, 13-3
- Total Seek Avg column, 11-3
- trace data, 15-6, 15-7
- trend data, 1-4
 - how CA-ASTEX uses, 1-5
- TSO commands, 13-2
- TTR program, 15-7

U

- Unbal% column, 2-4, 10-5
- unmovable data sets, 15-5
- UP command, A-2
- utilities
 - for reorganizing non-partitioned data sets, 15-7
 - for reorganizing partitioned data sets, 15-6
 - for reorganizing volumes, 15-2

V

- VIEW command, A-2
- Volser and Jobname, 6-2, 7-2, 8-2, 9-2, 11-2
- volume reorganization utility, 15-2